URLs SciDAC: www.osti.gov/scidac/
nal Collaboratories: DOEcollaboratory.pnl.gov
formance Networks: www.sc.doe.gov/ascr/mics/hpn

United Spaces for Distributed Science

Overview of National Collaboratories and SciDAC Network Research Projects

Mary Anne Scott, DOE
National Collaboratories Program

Thomas N
High-Perfo
Program

Ray Bair, PNNL

Scientific Discovery through Advanc
Annual Meeting
March 10-11, 2003
Napa, CA

OFFICE OF SCIENCE U.S. DEPARTMENT of ENERGY

Topics

- Collaboratory and Infrastructure Pilots
- Earth Systems Grid – Don Middleton
- Middleware Projects
- Network Research Projects

Collaboratory Pilots and Collaboratory Infrastructure

DOE Science Grid

National Fusion Collaboratory

Particle Physics Data Grid

Collaboratory for Multi-Scale Chemical Science

Earth System Grid

U.S. Department of Energy
National Collaboratories

The DOE Science Grid

Computing and Data Infrastructure For Large-Scale Science

William Johnston, PI, LBNL Ray Bair, co-PI, PNNL Ian Foster, co-PI, ANL Al Geist, co-PI, ORNL William Kramer, co-PI, NERSC	Engineering Working Group Keith Jackson, Chair, LBNL Tony Genovese, ESnet Mike Helm, ESnet Von Welch, ANL Steve Chan, NERSC Kasidit Chanchio, ORNL Scott Studham, PNNL

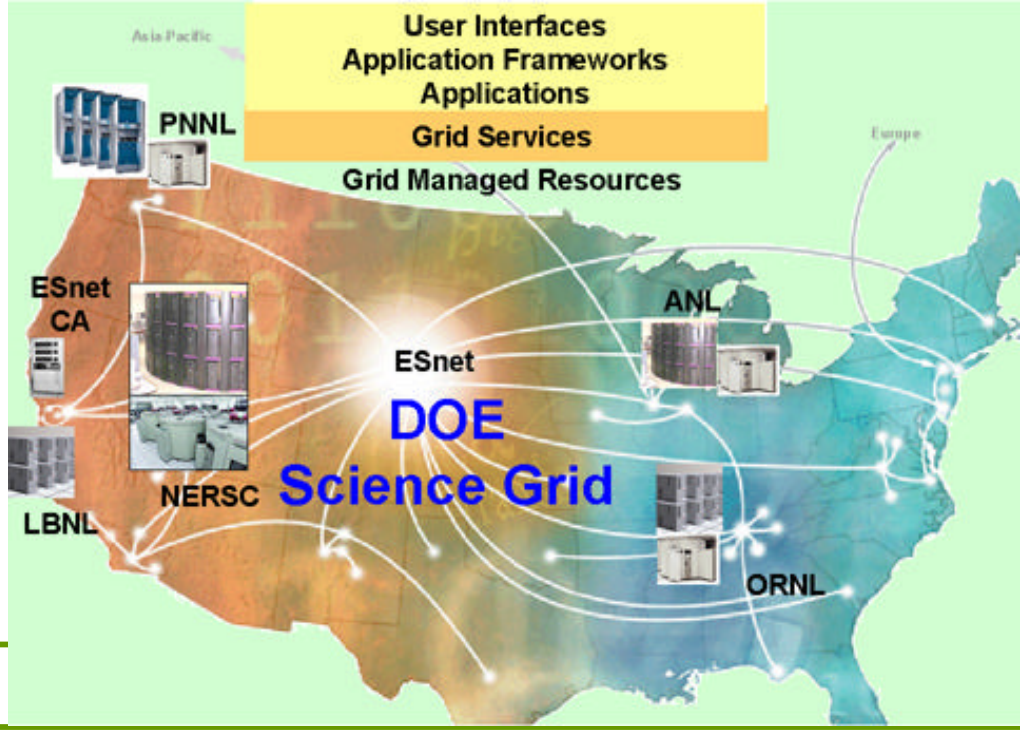

Objectives

- Grid technology dissemination to DOE science projects
- Middleware for uniform, secure, and highly capable access to large and small scale computing, data, and instrument systems, all of which are distributed across organizations
- Services supporting construction of application frameworks and science portals
- Persistent cyberinfrastructure for building and operating distributed applications (e.g. security services and resource access)
- Persistent resource infrastructure, e.g. NERSC on the Grid

Accomplishments

- Established effective inter-Lab coordination to build Science Grid and have resources committed to the Science Grid (e.g. PDSF at NERSC, Jazz cluster at ANL, Compact cluster at PNNL)
- Development of site firewall policy requirements for Grids
- Python wrapper for Globus – used for various monitoring applications
- Auditing and fault monitoring system
- Advances in Grid authorization - driven by HEP and Fusion
- ESnet DOEGrids CA project – a close collaboration w/ Science Grid
  - PKI authentication infrastructure supporting large-scale DOE science collaborations (HENP, Fusion, Science Grid)
  - Will sign subordinate CA certs (e.g. for Grid integration w/ Kerberos)
  - Provides a directory service for all certificates to support advanced authorization systems
  - Integrated policy with European Data Grid – HEP collaborators in US and Europe can now share resources, NERSC will accept EDG certs

Accomplishments (cont.)

- NERSC
  - Risk mitigation plan for NERSC systems to be on the Grid
  - NERSC funded HPSS integration with Grid was tested and validated on the Science Grid
  - IBM collaboration for Globus on AIX/SP
  - A unified Grid authorization system based on LDAP server across all the production hosts
  - NERSC accounting system integrated with Grid ids
  - Working toward using Grid Security Infrastructure and certificates as a foundation for a single sign-on at NERSC

Available Software and Technology

- **DOEGrids CA is available to the DOE-related science community to issue PKI identity certificates for collaborations**
- **pyGlobus (Python callable Globus modules)**
  - **NetSaint Grid status monitoring**
  - **Grid administration tools**
- **Grid helpdesk software**
- **Grid auditing and fault monitoring (prototype)**

Collaborations and Interactions

- Grid based Genome Analysis and Databases Update application (ANL)
- Authentication and authorization infrastructure for Particle Physics Data Grid (PPDG) and Fusion community
- Application “top-to-bottom, end-to-end” Grid performance monitoring for HEP/Atlas
- Cross-Grid resource sharing with iVDGL (HEP/Atlas)
- Grid testing platform for
  - ECCE chemistry workbench
  - PNNL subsurface transport model
- PDSF cluster at NERSC on Grid for Atlas and STAR HENP experiments

THE GOAL OF THE NFC IS TO ADVANCE SCIENTIFIC UNDERSTANDING & INNOVATION IN FUSION RESEARCH

Collaboratory is required to advance fusion science: geographically diverse community (37 states, 3 large experiments), leading to 1 worldwide experiment

- Diverse team
  - ANL: DSL & FL
  - GA: DIII–D Fusion Lab
  - LBNL: Distributed Systems
  - MIT: C–Mod Fusion Lab
  - Princeton Computer Science
  - PPPL: NSTX Fusion Lab
  - U. of Utah: Scientific Comp. & Imaging
- Objective is to advance fusion science
  - Experimental facilities
  - Integrate experiment, theory, modeling
  - Create a common toolkit for services

U.S. Department of Energy National Collaboratories

SciDAC Scientific Discovery through Advanced Computing

FusionGRID www.fusiongrid.org

FIRST YEAR ACCOMPLISHMENTS OF THE NFC

Grid Computing

- FusionGrid created: MDSplus data system integrated with Globus GSI
- FusionGrid released with complete monitoring: TRANSP & GS2 fusion codes remotely accessible via Globus/Akenti and fine-grain authorization via GRAM
- FusionGrid used for scientific calculations presented at the APS/DPP Mtg
- Prototyped: between pulse pre-emptive scheduling, parallel MDSplus I/O

Visualization

- SCIRun 3D visualization of NIMROD fusion data via MDSplus
- SCIRun visualizations used for scientific work presented at APS/DPP
- Access Grid functional on Tiled Wall as well as small scale system (PIG)
- Collaborative Visualization: Wall to wall/workstation (VNC, DMX), ELVis

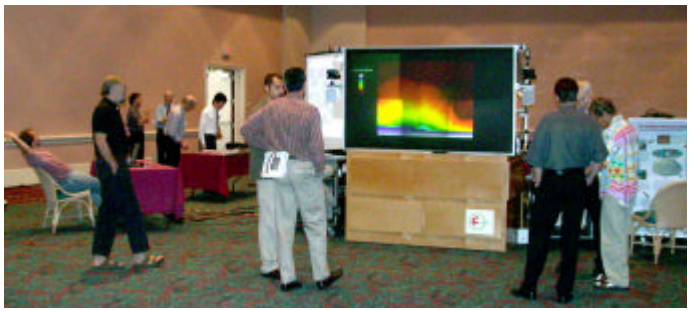

Culminated in a full demonstration at the Nov 2002 APS/DPP meeting attended by over 1000 fusion scientists

NFC'S TOOLS AND TECHNOLOGIES

- Secure MDSplus using Globus GSI available
  - Authentication and Authorization using DOE CA
- TRANSP available for worldwide usage on FusionGrid
  - Beowulf cluster, client application, complete job monitoring
  - Secure access by Globus GSI, Akenti, DOE Grids CA
- Personal Access Grid (PIG) software and spec's available
  - Installed at MIT and GA; PPPL has large AG node
- SCIRun for 3D visualization including MDSplus stored Fusion data
- Toolkits for sharing visualization wall to wall and on AG
  - Tiled walls at GA and PPPL

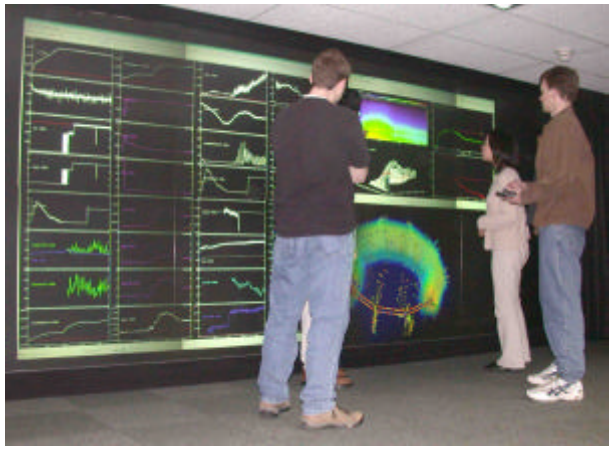

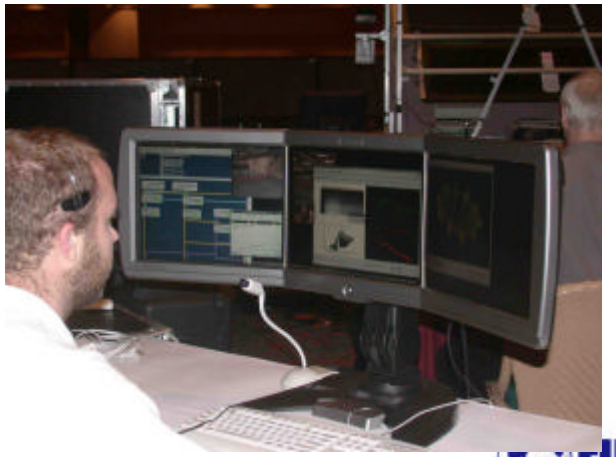

NFC'S COLLABORATIONS AND INTERACTIONS WITH OTHER PROJECTS

- SciDAC Center for Extended Magnetohydrodynamic Modeling
  - NIMROD data in MDSplus visualized by SCIRun presented at APS/DPP
- SciDAC Plasma Microturbulence Project
  - GS2 code being tested on FusionGrid for broader & easier usage
- Plasma Science Advanced Computing Institute (PSACI)
  - Strong endorsement of NFC plans & accomplishments by PSACI PAC
- Data Grid Toolkit; Security & Policy for Group Collaboration; Distributed Security Architecture
  - Secure access, authentication, authorization, Globus GSI/Akenti
- Particle Physics Data Grid; DOE Science Grid
  - Site security,Firewalls, and Grid security; CA for FusionGrid
- Middleware to Support Group to Group Collaboration
  - AG development: user education & testing & feedback
- eServices Infrastructure for Collab. Science; Portal Web Services
  - NFC & Fusion science as customer

Particle Physics Data Grid Collaboratory Pilot

Computer Science Groups: Condor, Globus, SRM, SRB

Nuclear Physics Experiments: STAR, TJNAF

Lattice QCD: TJNF

High Energy Physics: ATLAS, BaBar, CMS, D0

PI's: Richard Mount, SLAC, Miron Livny, Wisconsin, Harvey Newman, Caltech

Steering Committee: John Huth, Harvard (ATLAS), Tim Adye, RAL (BaBar), Lothar Bauerdick, FNAL (CMS), Lee Lueking FNAL (D0), Chip Watson, TJNAF, Jerome Lauret, BNL (STAR), Miron Livny, Wisconsin (Condor), Jennifer Schopf, ANL (Globus), Ian Foster, ANL (Globus), Reagan Moore, SDSC (SRB), Arie Shoshani, LBNL (SRM)
Coordinators: Ruth Pordes, FNAL, Doug Olson, LBNL

Liaisons: Paul Avery (iVDGL), Larry Price (HICB), Mike Wilde(GriPhyN), Torre Wenaus, Ian Bird (LCG)

Program Managers: Mary Anne Scott (MICS), Irwin Gaines (HENP), Steve Steadman (NP)

Experiment data handling requirements:

- Petabytes of storage, Teraops of computing, Thousands of users,
- Hundreds of institutions, 10+ years of analysis ahead

Focus of PPDG:

- Vertical Integration of Grid technologies into Applications ongoing work
- Deployment, hardening and extensions of common Grid services and standards – data replication, storage and job management, monitoring and planning.
- Interdisciplinary teams of physicists, engineers and computer scientists
- Using and extending grid technologies ranging across architecture, integration, deployment and robustness
- Driven by demanding end-to-end applications of experimental physics
- Enabling new scales of research in experimental physics and experimental computer science

U.S. Department of Energy National Collaboratories

SciDAC Scientific Discovery through Advanced Computing

PPDG

PPDG Accomplishments

Application "Grids" with incremental capabilities

- **Data Replication for BaBar** - Terabyte stores replicated from California to France and England.
- **Replication and Storage Management for STAR** and **JLAB** - development and deployment of standard API and interoperable implementations.
- Production Simulation Grids for **ATLAS** and **CMS, STAR** distributed analysis jobs.
- **Data Transfer, Job and Information Management for D0** - GridFTP integrated with SAM; Condor-G job scheduler, MDS resource discovery all integrated with SAM.

Initial Security Infrastructure for Virtual Organizations

- **PKI certificate management**, policies and trust relationships (using DOE Science Grid and Globus)
- Authorization mechanisms – standard callouts for Local Center Authorization for Globus, EDG,
- Prototyping secure credential stores
- Engagement of site security teams.

Data and Storage Management

- Robust data transfer over heterogeneous networks using standard protocols: GridFTP, bbcp
- Distributed Data Replica management: SRB, SAM, SRM
- Common Storage Management interface and services across diverse implementations: SRM - HPSS, Jasmine, Enstore

Job Planning, Execution and Monitoring

- Job scheduling based on resource discovery and status - condor-g and extensions
- Retry and Fault Tolerance in response to error conditions - hardened gram, gass-cache, ftsh,
- Distributed monitoring infrastructure for resource discovery, resource and job information - MDS, Monalisa, Hawkeye

Prototypes and Evaluations

- Grid enabled physics analysis tools
- End to end troubleshooting and fault handling
- Cooperative Monitoring of Grid, Fabric, Applications

Technologies, Tools and Applications

Who uses what?

	Atlas	BaBar	CMS	D0	STAR	TJNAF	TJNAF-LQCD
Globus GridFTP	x	bbcp	x	x	x	x	
Globus MDS	x		x	x			
Globus GSI	x	x	x (+KCA)	x (+KCA)	x	x	Java
Condor-G/GRAM	x/VDT	x/EDG	x/VDT	x			
SRM			x	x	x	x	
SRB		x	x				
Distributed Data Mgt.	MAGDA	Bbserver++	MOP	SAM	x	x	x
Analysis prototypes		JAS	Caigee Clarens	root/sam			
Automated Data Replication	x	x		x	x	x	x
Applications Grid	p	p/EDG	x	x/sam			
DOE SG Certs	x	x	x	x	x	x	x

PPDG Collaborations and Interactions

US Physics Grid Projects

GriPhyN, iVDGL, PPDG – the Trillium + Virtual Data Toolkit

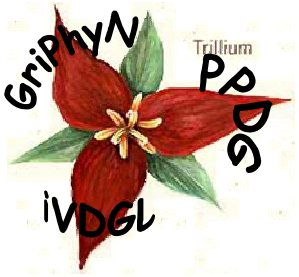

EU-US Physics Grids

LHC Computing Grid – ATLAS, CMS

European Data Grid - BaBar, D0, ATLAS, CMS

High Energy Physics Intergrid Coordination Board

SciDAC projects	**Interaction**
Earth System Grid II	SRM
DOE Science Grid	CA, RA
A High-Performance Data Grid Toolkit	Globus Toolkit
Storage Resource Management for Data Grid Applications	SRM
Security and Policy for Group Collaboration	GSI, CAS
Scientific Data Management Center	STAR SDM
Bandwidth Estimation: Measurement Methodologies and Applications	IEPM-BW
A National Computational Infrastructure for Lattice Gauge Theory	TJNAF/LQCD
Distributed Monitoring Framework	Netlogger, Glue Schema

U.S. Department of Energy National Collaboratories

PPDG

Collaboratory for Multi-scale Chemical Sciences (CMCS)

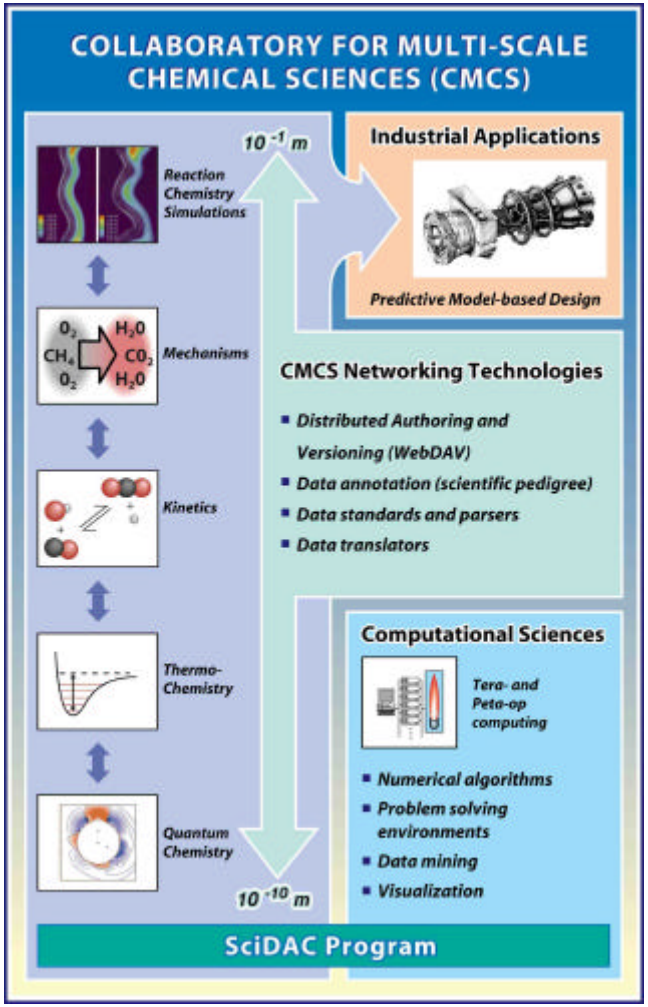

- A collaboration of eight national labs and universities
- Objectives
  - Architect and build an adaptive informatics infrastructure enabling multi-scale science
    - XML Data/metadata services
    - Chemical Sciences Portal
  - Pilot project within combustion community
    - Rapid exchange of multi-scale data/pedigree
    - Integrate Chemical Science applications
  - Demonstrate the power of adaptive infra-structure as CMCS evolves, and to new areas
  - Gain adoption/continued support by science community participation
  - Document success & continuation path

MICS Mathematical, Information, and Computational Sciences · Sandia National Laboratories · Pacific Northwest National Laboratory Operated by Battelle for the U.S. Department of Energy · Argonne National Laboratory · Los Alamos National Laboratory · NIST National Institute of Standards and Technology · MIT Massachusetts Institute of Technology · Berkeley University of California

U.S. Department of Energy National Collaboratories

CMCS Prototype Demonstrates Portal, Data Pedigree, Application Integration

CMCS Explorer Portlet provides rich interface to multi-scale data/metadata and services that include automatic data translations.

CMCS Portal environment, a CHEF Jetspeed collaboration

Thermochemical Active Tables (ATcT) and webservice a CoG Kit collaboration.

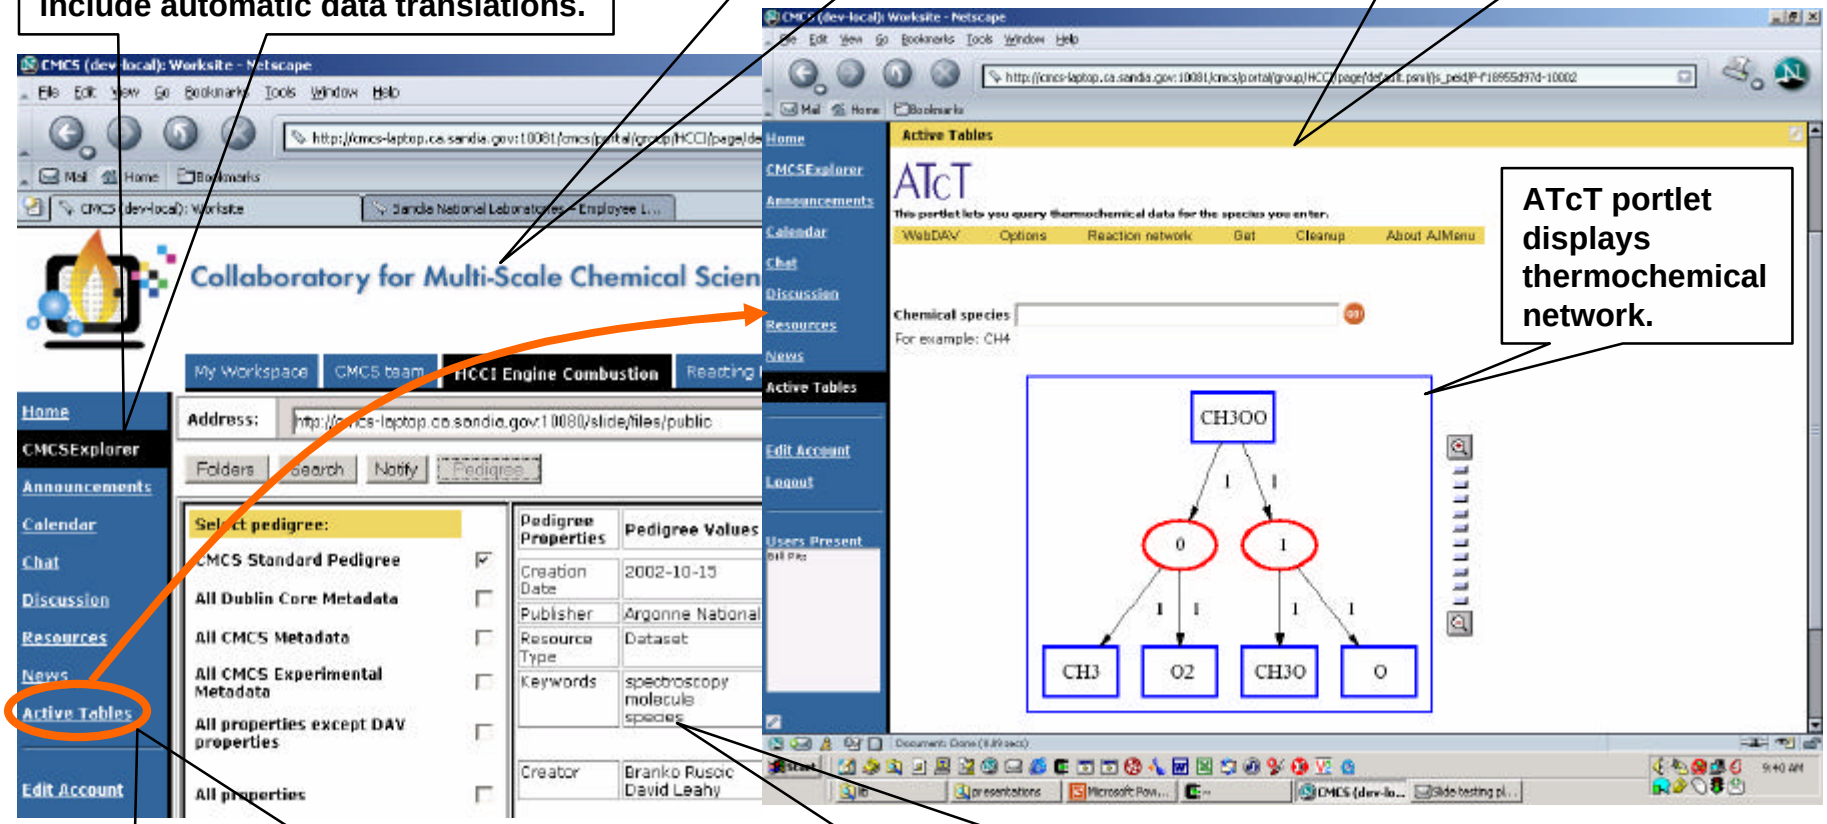

ATcT portlet displays thermochemical network.

ATcT Web service accessed from portlet in the HCCI Project Team workspace in the CMCS Portal

Pedigree of data viewed from CMCS Explorer with links and annotation enables 'Pedigree Browsing'.

CMCS Tools and Technologies

- CMCS Data/Metadata services
  - SAM translation, annotation
  - WebDAV implementation
  - Notification (JMS, NED)
  - Search
  - Pedigree browsing
  - Core XML schema
- Chemical Science Portal
  - Jetspeed (CHEF)
  - Explorer and application portlets
  - Community services
- Application Integration
  - Webservices
  - WebDAV API
  - Multi-scale data including NIST access

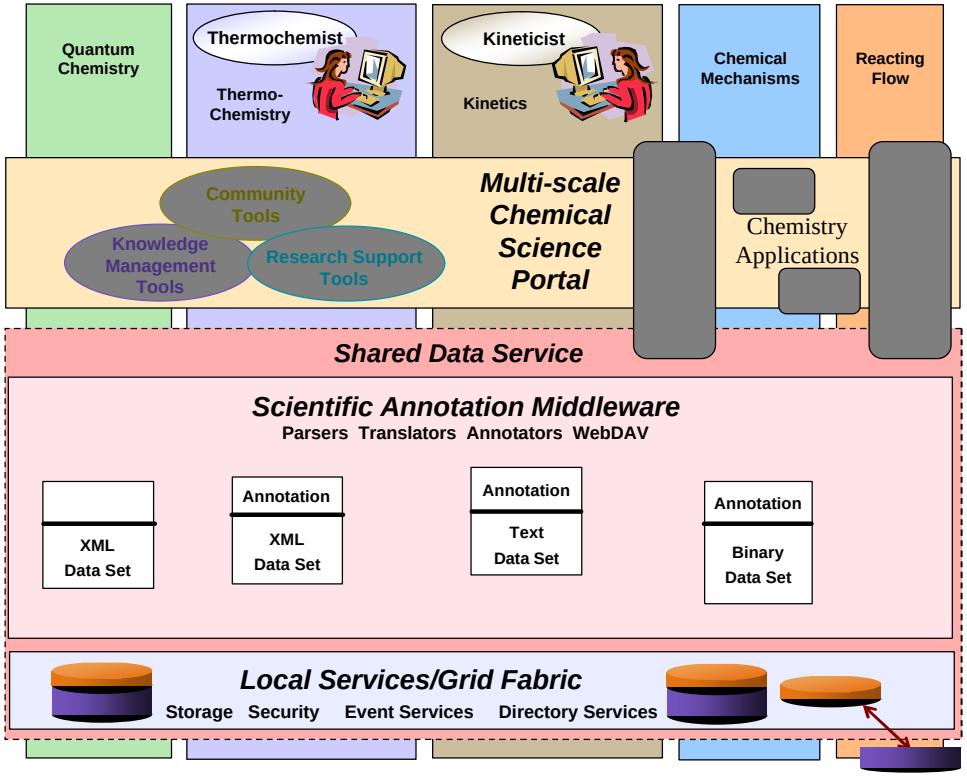

A diagram representing the major conceptual elements of the CMCS Informatics Infrastructure.

CMCS Collaborations & Interactions

- Scientific Annotation Middleware Project
- CoG Kit project
- CHEF Jetspeed
- CCA ISIC
- SDM ISIC
- BES Chemical Science SciDAC projects
- Earth Systems Grid

CMCS Team

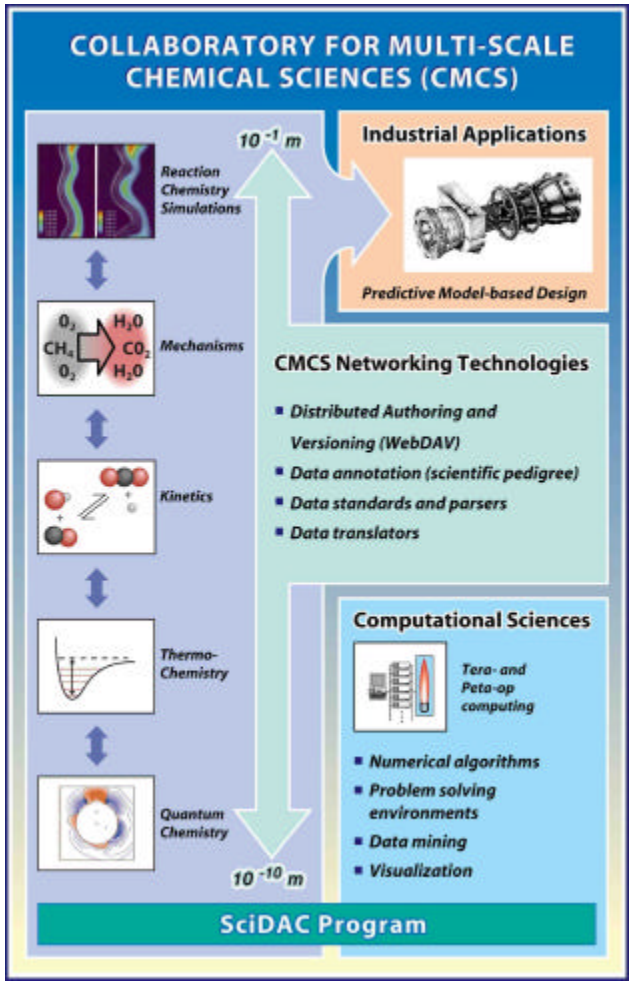

- Larry Rahn-SNL*, Christine Yang, Carmen Pancerella, Wendy Koegler, David Leahy, Michael Lee, Renata McCoy,
- Theresa Windus-PNL*, James D. Myers, Karen Schuchardt, Brett Didier, Eric Stephan, Carina Lansing,
- Al Wagner-ANL*, Branko Ruscic, Michael Minkoff, Sandra Bittner, Gregor von Laszewski, Reinhardt Pinzon, Sandeep Nijsure, Kaizar Amin, Baoshan Wang,
- William Pitz-LLNL*,
- David R. Montoya-LANL*, Lili Xu, Yen-Ling Ho,
- Thomas C. Allison-NIST*,
- William H. Green, Jr.-MIT*,
- Michael Frenklach-UCB*

* denotes Institutional Point of Contact

SAM

The Globus Project™ Java CoG Kit www.globus.org/cog

Xpowered by Chef Jetspeed POWERED

Office of Basic Energy Sciences — Serving the Present, Shaping the Future

MICS Mathematical, Information, and Computational Sciences

Sandia National Laboratories

Pacific Northwest National Laboratory — Operated by Battelle for the U.S. Department of Energy

Argonne National Laboratory

Los Alamos NATIONAL LABORATORY

NIST National Institute of Standards and Technology

MIT Massachusetts Institute of Technology

Berkeley University of California

U.S. Department of Energy National Collaboratories

The Earth System Grid

http://www.earthsystemgrid.org

- *Enabling the simulation and data management team in organizing and managing terascale data*
- *Enabling the research community in analyzing and visualizing results*
- *Enabling broad multidisciplinary communities to easily discover and access simulation results*

PARTICIPANTS

ANL – Ian Foster (PI), Veronika Nefedova, (Bill Allcock), (John Bresenhan), (Joe Link)

LBNL – Arie Shoshani, Alex Sim

LLNL/PCMDI – Bob Drach, Dean Williams (PI)

NCAR – David Brown, Peter Fox, Jose Garcia, Don Middleton (PI), Gary Strand, Luca Cinquini, (Lawrence Buja)

ORNL – Dave Bernholdt, Kasidit Chancio, Line Pouchard

USC/ISI – Ann Chervenak, Carl Kesselman, (Laura Perlman)

Accomplishments

At SC2002 we demonstrated a Prototype ESG Web Portal for climate data that integrated a number of distributed technologies/services:

- Security: authentication and authorization
- A Metadata Catalog with search capability
- Dataset replica management
- Data transport/management capabilities including interoperability among NCAR MSS and DOE HPSS systems
- Server for new OPeNDAPg data access protocol built upon Striped GridFTP server (alpha)
- Web-based browsing and visualization
- Additional work is underway now: client side for OPeNDAPg, new metadata services architecture, new authorization structure, & full multi-user version of data management/transport tools

Tools & Technologies

- Winter 2002/2003
  - First user evaluation of data management & transport tools
- Spring 2003
  - New metadata schemas and catalog services
  - First user evaluation of web-based data search
  - First client analysis applications running over OPeNDAPg
- Summer 2003
  - First ESG Portal user evaluation - including security, search, browse, access/transport, visualize, status monitoring

Collaborations & Relationships

- CCSM Data Management Group
- The Globus Project
- Other SciDAC Projects
  - Climate, Security & Policy for Group Collaboration
  - Scientific Data Management ISIC
  - High-performance DataGrid Toolkit
- OPeNDAP/DODS (multi-agency)
- NSF National Science Digital Libraries Program (UCAR & Unidata THREDDS Project)
- U.K. e-Science and British Atmospheric Data Center
- Earth Science Portal group (multi-agency, int'l.)

Middleware for Portals and Spaces

Mary Anne Scott, DOE
National Collaboratories Program

Middleware to Support Group to Group Collaboration

Middleware Technology to Support Science Portals

Portal Web Services

Pervasive Collaborative Computing Environment

U.S. Department of Energy
National Collaboratories

Middleware to Support Group to Group Collaboration

- PI: Rick Stevens – ANL
- Co-PI: Terry Disz, Mark Hereld, Ivan Judson, Robert Olson, Michael E. Papka
- Goals:
  - Scalable peer to peer Virtual Venue server for the AG
  - Development and integration of shared workspaces into the AG
  - Improving and simplifying AG node management
  - Integration of security into AG infrastructure
  - Better integration of asynchronous collaboration tools

Accomplishments

- Design and Implementation of Access Grid 2.0
  - Produced of both design and architecture documents for review by public (beginning introduction into GGF document process)
  - Demonstrated full-featured prototypes at SC2002 of new venue architecture, venue client, workspace docking complete with application sharing
  - Improved upon prototype, for release version of Access Grid 2.0 (currently available in alpha release)
- Planning Access Grid Retreat 2003 with an emphasize on greater community involvement.

Available Tools and Technology

- Access Grid 2.0 Toolkit
  - Alpha (available now)
    - www.mcs.anl.gov/fl/research/accessgrid/software
  - Beta (March 15)
  - Final (April 15)
- Release includes
  - Scalable virtual venue server
  - Virtual venue server management tool
  - Node management tool
  - Virtual venue client
  - Integration of Globus based security

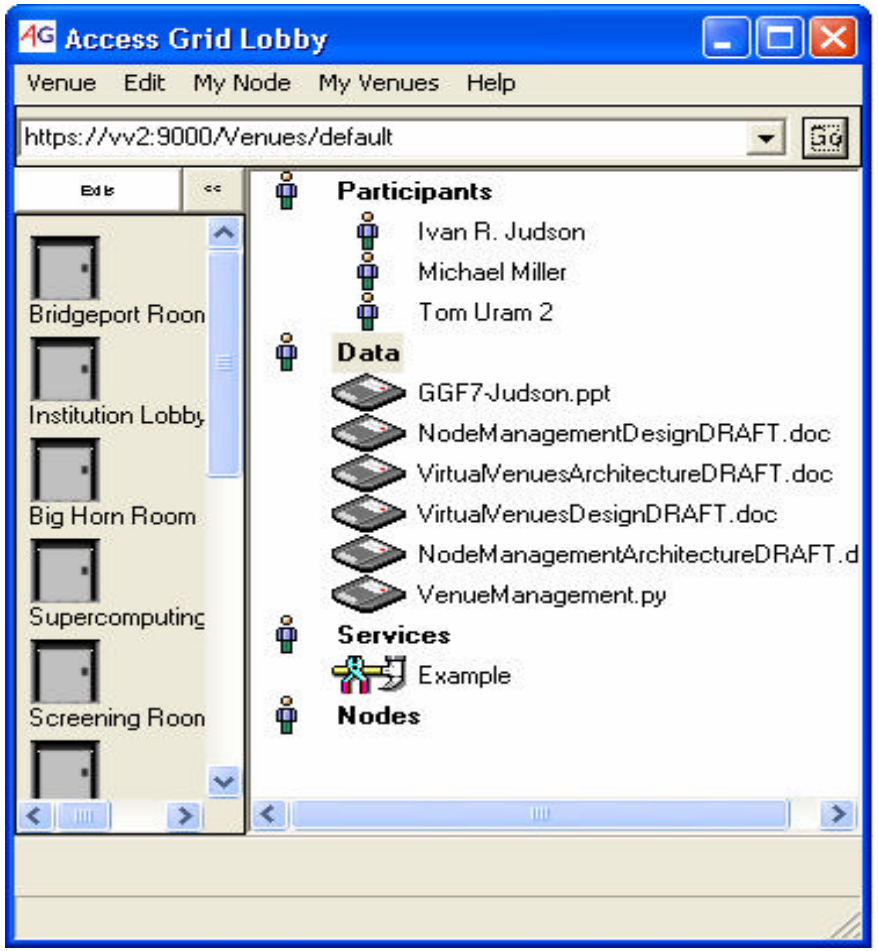

Access Grid v2.0 Venue Client

Collaborations and Interactions

- The National Fusion Grid
  - Introduction of AG to Fusion community
  - Major factor in the development of Personal Interface to the Grid (PIG) (desktop Access Grid solution)
- Reliable and Secure Group Communication
  - Integration of InterGroup Chat as mud replacement
- Access Grid Support
  - Earth Science Grid
  - CoG Kits: Enabling Middleware for Designing Science Applications, Web Portals and Problem Solving Environments

Middleware Technology to Support Science Portals: a Gateway to the Grid

Dennis Gannon, Randall Bramley
Dept of Computer Science, Indiana University

- Project Objectives
  - To develop the middleware to allow users easy access to Grid resources and services
  - A Web portal to support the work of collaborating groups of scientists who need access to Grid technology.

Project Accomplishments

- Original prototype used in US Atlas “Grappa” portal
- Defined a new reference architecture
  - In collaboration with Grid Forum “Grid Computing Environment” working group.
  - Uses Java “Portlet” technology
    - As supported in Apache Jetspeed and IBM Websphere
  - Extended to support interaction with Globus and Grid web services
- New prototype in use in NCSA Grid portal
  - Also components now incorporated into NEES Grid portal

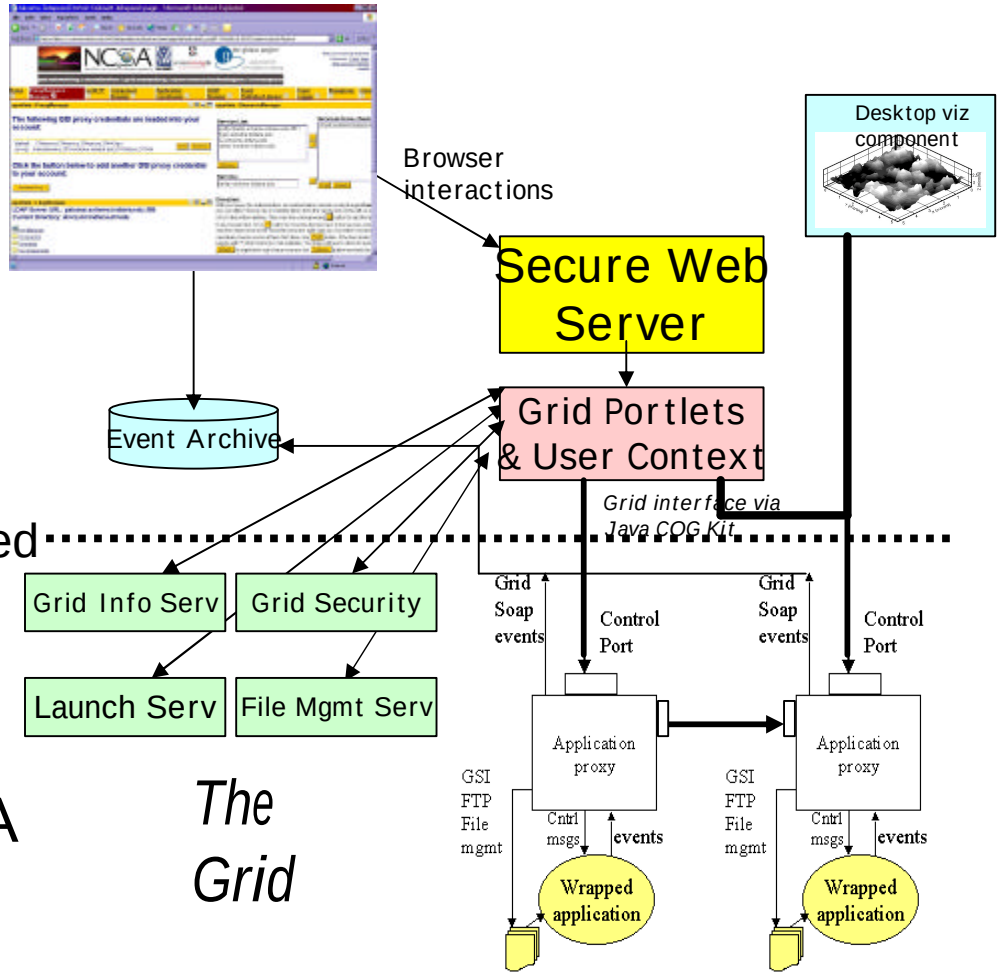

See http://www.extreme.indiana.edu/xportlets

Available Tools and Technology

Provided Capability

- Management of user proxy certificates
- Remote file Management via Grid FTP
- News/Message systems
  - for collaborations
- Event/Logging service
- Access to OGSA services planned
- Access to directory services
- Specialized Application Factories
  - Distributed applications
  - Workflow (in progress)
- Access to Metadata Index tools
  - (in progress)

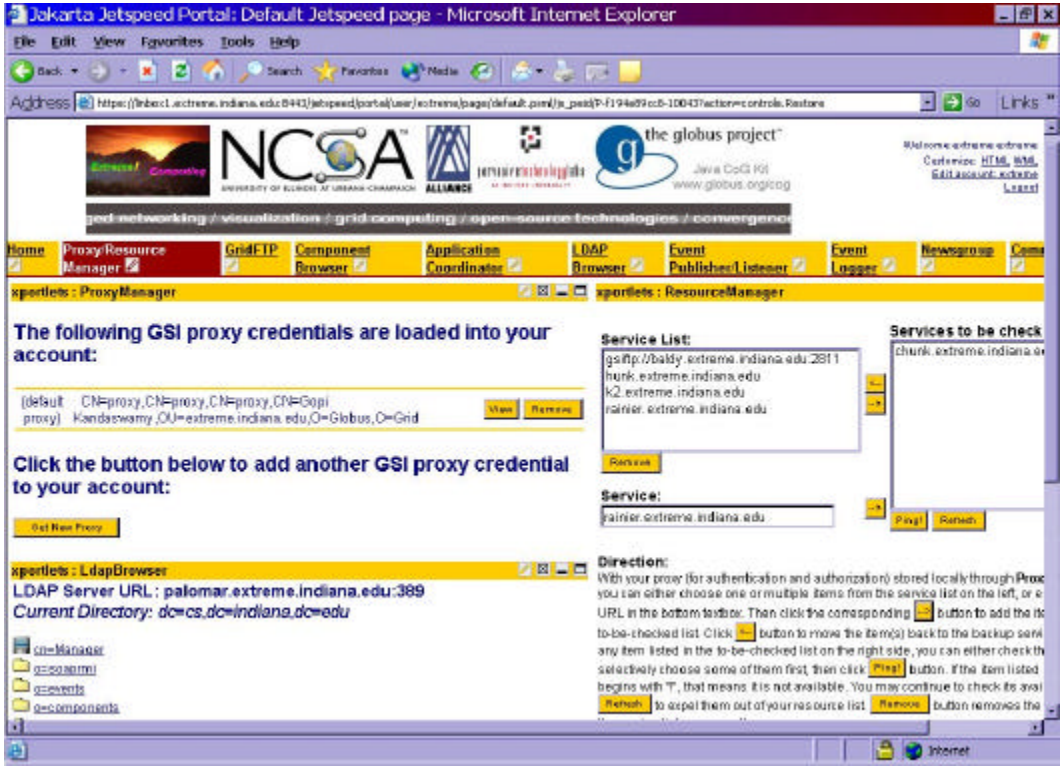

See
http://www.extreme.indiana.edu/xportlets

Collaborations

- Grid Forum
  - Grid Computing Environments WG
  - OGSA WG
  - Portal Architecture workshop at GGF7
- NCSA Portal Expedition
  - With Indiana Pervasive Tech labs & Argonne & Utah
- NASA IPG
- SciDAC CCTTSS
- Grid Lab (EU)
- NPACI via Univ. of Texas
- NEES Grid via Univ of Michigan

Portal Web Services: Support of DOE SciDAC Collaboratories

- Mary Thomas (U. Texas), Geoffrey Fox, Marlon Pierce (Pervasive Computing Lab, Indiana), Reagan Moore (UCSD), Dave Schissel (General Atomics)
- Project Objectives
  - Develop generalized toolkits for building portal and grid web services specialized for DOE environment
  - Deploy portal and web services based on toolkits across multiple SciDAC Collaboratories – initially Fusion Grid
  - Technology transfer of toolkits and experiences to the DOE user community.

Project Accomplishments

- Proposed SciDAC Fusion Architecture
  - Grid Information Archival web service
  - GridPort 2.0 auto install package ready to go
- SRB
  - MDS+ integration planned
  - WSDL/Web service, OGSA version planned
- Developing Portal Component Architecture (Portlets/OGSA) with GCE
- Jetspeed Portal (in collaboration with Indiana)
  - Portlets based on Perv. Techn. Labs, TACC, IU, U. Mich, Alliance
  - Handle remote sessions and load remote websites.file transfer, interportlet communication, newsgroup system to supports email and browser based access; planning SRB, LSF, Grid-IAS
  - Interface to GridPort toolkit

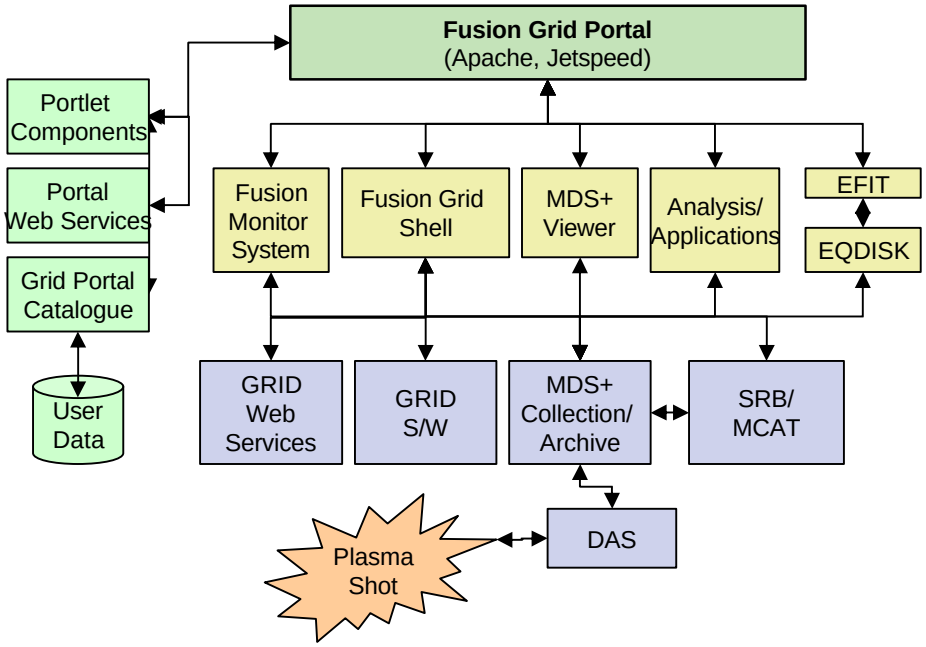

Available Tools and Technologies

- Grid Information Archive Service
  - client and web service available
  - Full information web service; VO based
- GridPort Toolkit:
  - GridPort 2.0 released: NMI R2 compliant; auto-install, SRB 2.0; runs on Apache, Netscape, Sun, Linux
- SRB 2.0 released
- Integration portals at TACC and PCL

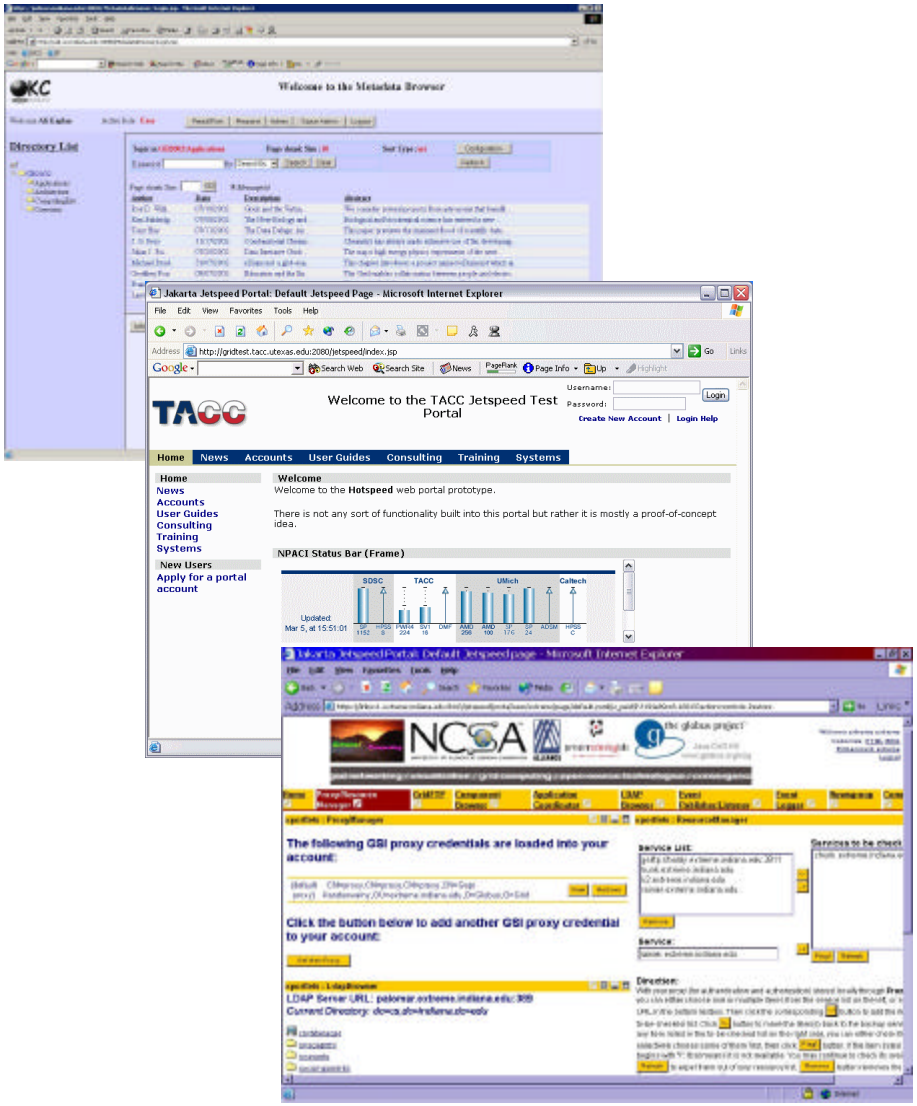

Collaborations

- PACI and TeraGrid:
  - NPACI -SDSC, U. Mich (NEESGrid);
  - Alliance: NCSA, Indiana;
- Globus: Argonne, ISI
- Global Grid Forum: Grid Computing Environments; Persistent Data Archives
- Gov: NASA IPG, DoD PET
- Texas Internet Grid for Research and Education (TIGRE): U. Texas, U. Houston, Texas A&M, Texas Tech, Rice
- Commercial: Platform Computing, United Devices, IBM

FusionGRID www.fusiongrid.org

TACC

SDSC

NSF

U.S. Department of Energy National Collaboratories

pervasive technology labs AT INDIANA UNIVERSITY

NPACI

ALLIANCE

Pervasive Collaborative Computing Environment (PCCE)

LBNL Personnel: Deb Agarwal, Charles McParland, and Marcia Perry

UW Personnel: Miron Livny

Objective: Support the day-to-day collaboration needs of scientists by building tools that leverage off the web, Grid services, existing standards, and existing tools. The final environment needs to support a continuum of collaboration capabilities.

Accomplishments/Highlights

- Developing a collaborative workflow tool
  - Web Services Flow Language-based prototype workflow engine developed
  - Specified Condor and CoG job submission mechanisms
- Designed and implemented a prototype presence and messaging tool which provides the baseline environment for users
- Working with Global Accelerator Network and Large Hadron Collider experiments and accelerator physicists to identify and understand collaboratory tool needs
- Published a paper describing the PCCE

Available Tools and Technologies

- Secure messaging and presence tool
  - Open source
  - Based on modified IRC
  - Graduated security including Grid certificates
  - Beta release April 2003 (alpha release in progress)
- Reliable file transfer – developed with CoG project
- Collaborative workflow tool
  - Still under development
- Publications describing PCCE
- URL: http://www-itg.lbl.gov/Collaboratories/pcce.html

Connections

- ANL – Access Grid
- CMS and Atlas collaborators
- Global Accelerator Network (GAN)
- Distributed security architectures project
- Scalable and secure peer-to-peer information sharing project
- Reliable and secure group communication
- Advanced collaborative environments working group of Global Grid Forum

Middleware for Data Transport and Management

High-Performance Data Grid Toolkit

Storage Resource Management for Data Grid Applications

Scientific Annotation Middleware

the globus project™
www.globus.org

High Performance Data Grid Toolkit

- Co-PIs: Ian Foster (ANL), Carl Kesselman (USC/ISI), Miron Livny (U. Wis)
- Other Senior Personnel:
  - Bill Allcock (ANL), John Bent (U. Wis), John Bresnahan (ANL), Ann Chervenak (USC/ISI), Joe Link (ANL), Bob Schwartzkopf (USC/ISI), Doug Thain (UWis), Steve Tuecke (ANL)
- Objective: efficient, high performance, reliable, secure, and policy aware management of large-scale data movement.

the globus project™
www.globus.org

Project Accomplishments

- Achieved 2.8 Gbs transfer rate with GridFTP
- Transferred 236 GB in 54 hours surviving multiple failures with no human intervention
- Our tools are in use in multiple production environments and helping produce real results
- Redesigned replica tools to significantly increase performance and scalability
- Multiple standards documents in front of the Grid Forum for standardization

Available Tools and Technologies

- Currently Available:
  - Non-striped GridFTP, 1st Generation replica tools
- In Alpha:
  - 2nd Generation replica tools, OGSA compliant Reliable File Transfer (RFT) service, storage management tools (NeST, Kangaroo, Data Placement (DaP) scheduler) and the fault tolerant shell
- Future:
  - Striped GridFTP, non-TCP transports, improved reservations and scheduling

Collaboration and Interaction

- Multiple SciDAC projects consider our software an integral part of their project:
  - ESG, PPDG, DOE Science Grid, Fusion (planned)
- We work closely to ensure that we incorporate the latest security advances:
  - CAS integration is ongoing in our tools.
- Many other projects, both US, and international are also extensively using our tools:
  - GriPhyn, iVDGL, EU DataGrid, NorduGrid (Scandinavia), GridPP (UK), NPACI, the Alliance

Storage Resource Management for Data Grid Applications

- Project participants
  - LBNL: Arie Shoshani (PI), Alex Sim, Junmin Gu
  - Fermilab: Don Pertavick (Co-PI), Timur Perelmutov
- Project objectives
  - Develop concepts, functionality & software for *managing storage resources* on the Grid
  - Complement compute and network resource management
  - Interoperate with other Grid Middleware components
- Storage resource Managers:
  - Provide shared storage resource allocation & scheduling
  - Especially important for data intensive applications
  - Manage grid access to files that are archived on a mass storage system (MSS)
  - Maximize file reuse => minimize file transfers
  - Automate management of thousands of files through a single request
  - automatic "garbage collection" of files on grid resources

Project Accomplishment: Robust File Replication Using SRMs

- Robust File replication of thousands of files using a single command
  - Queues file request, reuses space automatically
  - Recovers from transient transfer failures
    - HPSS staging and archiving
    - Network failures
  - Used in production BNL – LBNL
- SRMs Used regularly at
  - LBNL (ESG + PPDG)
  - BNL (PPDG)
  - NCAR (ESG)
  - ORNL (ESG)
  - Fermilab (PPDG)

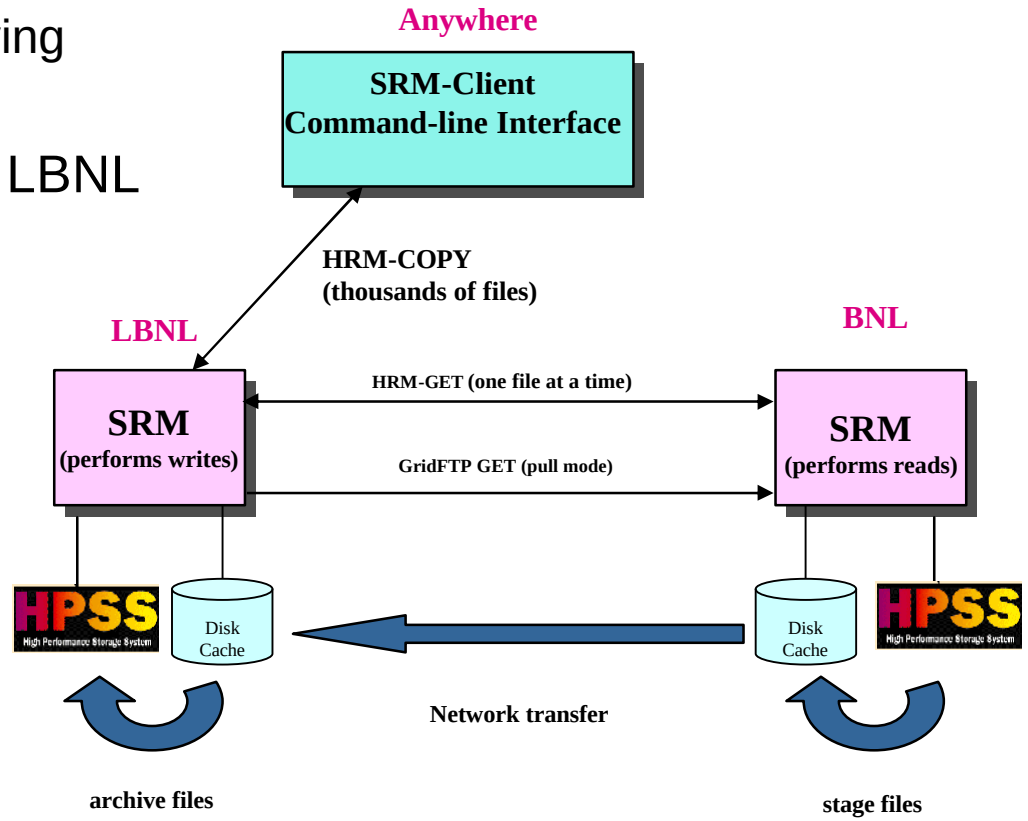

Available Tools and Technologies

- Technology
  - Design documents, papers, user guides
  - Interface specification – in CORBA and Web Service (WSDL)
  - Coordination of a standard with US - PPDG and European Data Grid participants
    - To foster independent interoperable implementations
  - Demonstrated interoperation between LBNL, JLAB, Fermilab, CERN
    - Each having a different Mass Storage System
- Software
  - A Disk Resource Manager (DRM) and a Hierarchical Resource Manager (HRM) for mass storage systems
  - An HRM for HPSS
    - Installed at LBNL, BNL, ORNL
  - A HRM for NCAR's MSS
    - Installed at NCAR
  - A HRM for Enstore
    - Installed at Fermilab
  - A Web-based File Monitoring Tool of file transfers

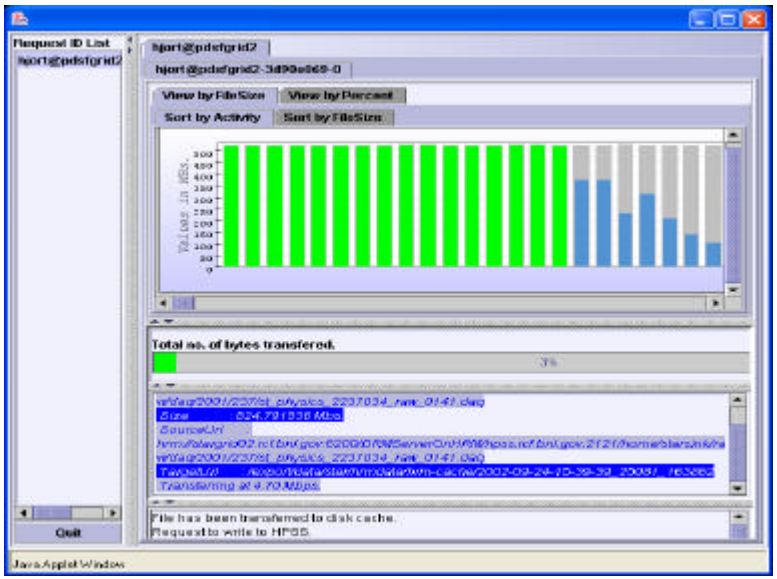

http://sdm.lbl.gov/srm

collaborations and interactions with other projects

- Particle Physics Data Grid
  - File replication in the STAR experiment
  - Grid multi-file access from multiple locations
  - Grid access from HPSS
  - Interoperability with CLAS experiment (Jlab)
  - Interoperability with Daisy experiment (Fermilab)
- Earth Systems Grid
  - Grid multi-file access from multiple locations
  - Grid access from NCAR's MSS
- Scientific Data Management (SDM) ISIC
  - Used for developing dynamic file access for HENP analysis
  - Used for developing co-scheduling algorithms of compute and storage resource
- Coordination with European Data Grid
  - Lead joint design
  - Being applied to MSS, (called Castor) at CERN
  - Being considered at Ratheford Lab in England

Scientific Annotation Middleware (SAM)

Jim Myers, Alan Chappell, Matt Elder, Mike Peterson from PNNL

Al Geist, Jens Schwidder, David Jung from ORNL

Objectives

- Develop a lightweight, flexible middleware to support the creation and use of metadata and annotations
  - owner, pedigree, classification, dependencies, …
- Enable the sharing of annotations among portals and problem solving environments, software agents, scientific applications, and electronic notebooks
- Improve the completeness, accuracy, and availability of the scientific record.

Highlight: SAM based CMCS Pedigree Browser Allows Chemists to Browse Data and Resources

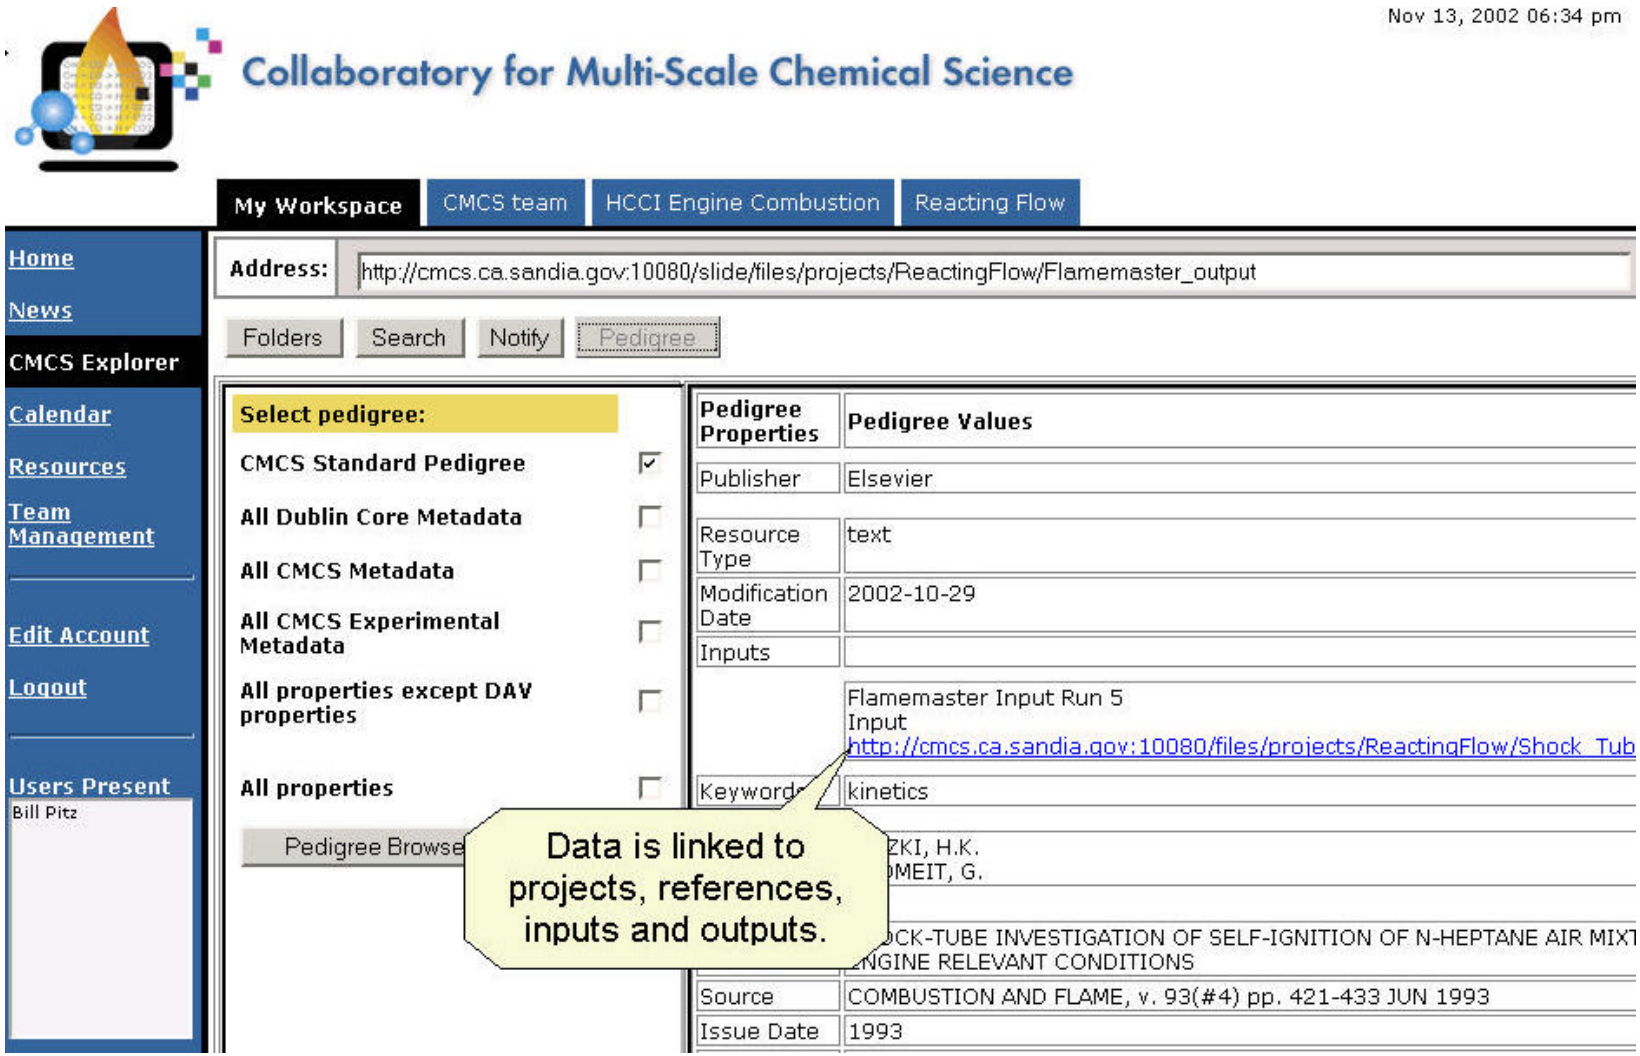

Available technologies
SAM Metadata Services Layer

- Jakarta Slide DAV server plus configurable:
  - Property Generation from binary/ASCII/XML files
  - Resource Translation
  - Mapping to Data Store(s)
  - JMS Events for access and changes
  - Authentication/Authorization model

Interactions with other projects

The DOE Collaboratory for Multiscale Chemical Science (CMCS) is using SAM as a component of a portal-centric system designed to facilitate collaboration, data exchange, and provenance tracking

CMCS is also using SAM to link the webDAV-aware Extensible Computational Chemistry Environment and chemistry applications.

The DOE Genomes to Life program is leveraging SAM to create a "biology aware" electronic notebook that can store and manipulate biological data objects.

U. Michigan CHEF project, Java CoG, Jakarta Slide, JCP JSR170, Grid Data Format and Information Retrieval Groups, Columbia U.

Research sponsored by Mathematics, Information and Computational Sciences Office
U.S. Department of Energy

Middleware for Secure Sharing, Multicast and Authentication

A Scalable and Secure Peer-to-Peer Information Sharing Tool

Reliable and Secure Group Communication

Distributed Security Architectures

U.S. Department of Energy
National Collaboratories

A Scalable and Secure Peer-to-Peer Information Sharing Tool

Personnel

K. Berket, D. Agarwal,
O. Chevassut, A. Essiari,
A. Muratas, and M. Thompson

Lawrence Berkeley National Laboratory

Objective: Create a peer-to-peer system to support location independent information sharing in the scientific community

Accomplishments/Highlights

- Created a prototype implementation of RDMF (Resource Discovery Messaging Framework), a simple messaging framework for resource discovery (based on XML)
- Created a prototype implementation of discofish, a concrete instantiation of RDMF
  - Enabled information discovery between two hosts using HTTP service for communication
- Created prototype of information sharing application (scishare) that allows users to:
  - Search for files by specifying a portion of a file name
  - Download a file whose metadata was returned as a result of a search
  - View the status of the transfer process for files that have been selected for download

Available Tools and Technologies

- Initial alpha release scheduled for Q2 2003
- Software in Java (requires JDK 1.3 or higher)
  - RDMF – messaging framework for resource discovery
  - discofish – concrete implementation of RDMF (with InterGroup support)
  - scishare – information sharing application with GUI front-end (with local information management)

Connections

- Reliable and Secure Group Communication
- Distributed Security Architectures
- Distributed Monitoring Framework
- Pervasive Collaborative Computing Environment

Reliable and Secure Group Communication

Personnel:

- Deb Agarwal, Karlo Berket, Guillaume Egles, and Olivier Chevassut at Lawrence Berkeley National Laboratory

Objectives:

- Developing the InterGroup protocol which is a reliable multicast protocol intended to scale to large group sizes and to the Internet.
- Developing a Secure Group Layer that provides a secure channel abstraction similar to Secure Sockets Layer (SSL), but for multicast

Accomplishments/Highlights

- InterGroup reliable multicast
  - Development of Java implementation
  - Implementation of Java and C++ interfaces
  - Development of a shared library of functions to make Java to C++ porting easier
- Secure Group Layer (SGL)
  - Developed and published required cryptographic algorithms and proofs of security
  - Designed framework for SGL implementation
  - Prototype SGL implementation (pre-alpha)

Tools/Technologies

- InterGroup implementation
  - Java and C++ interfaces
  - Demonstration chat application (used in AG 2.0)
- Secure Group Layer
  - Publications describing cryptographic algorithms and proofs of security for group key agreement
  - Publication describing SGL
  - Implementation of SGLv2 under way

Connections

- ANL – Access Grid
- PCCE Collaborative tools
- Scalable and secure peer-to-peer information sharing
- Distributed security architectures
- IETF – Reliable Multicast Transport Group
- IRTF – Reliable Multicast Research Group

Distributed Security Architectures

People: Mary Thompson, Abdelilah Essiari, Keith Beattie

Lawrence Berkeley National Laboratory

Goal: provide useable policy-based access control for computer mediated resources in distributed environments, such as collaboratories or computational Grids.

Our authorization service is currently compatible with Globus Toolkit 2, and we intend to modify our interfaces to conform to the emerging OGSA specifications.

Akenti Authorization Service

- PKI based
  - X.509 certificates identify users
  - digitally signed certificates to hold policy and authorization assertions
- Targeted at distributed environments
  - users, resources, stakeholders are geographically and administratively distributed
- Ease of use
  - Provides GUI and command line interfaces to create policy certificates
  - Provides simple API to check access, either as a library or a server

Accomplishments to Date

- Released code is available for
  - Akenti authorization library with C++ and C interfaces
  - Akenti standalone authorization server
  - support for runtime conditions
  - GUI and command line tools to create and verify policies
  - ***http://www-itg.lbl.gov/Akenti/download.html***
- Globus integration
  - Supports delegated proxy certificates
  - integrated with the Globus job-manager to do access control based on job definition parameters (RSL)

Interactions with other projects

- National Fusion Collaboratory
  - Experiment with policies to enforce priorities and fair use of resources
  - Called from legacy applications
- Secure and Reliable Group communications
  - Experiment with group policies and pre-authorization
- DOE Science Grid
  - Run Akenti authorization server as a web service for various other projects to experiment with.

Middleware for Grid Services and Instrumentation

Commodity Grid Kits (CoGkits)

eServices Infrastructure for Collaborative Services

Distributed Monitoring Framework Project

SciDAC Commodity Grid Kits (CoG Kits)

- Project Coordination
  - Gregor von Laszewski (ANL)
  - Keith Jackson (LBNL)
  - 3 Year project
- Objectives
  - Allow application developers to make use of Grid services from higher-level frameworks such as Java and Python.
  - Easier development of advanced Grid services.
  - Easier and more rapid application development.
  - Encourage code reuse, and avoid duplication of effort.
  - Encourage the reuse of Web Services as part of the Grids (now part of a separate proposal).

Sample Project Accomplishment Highlights

- Climate Modeling with the ANL Espresso Scientific Grid Interface
  - DOE Argonne Regional Climate Center (John Taylor)
  - Performed many multi-year high resolution regional climate simulations using a version of MM5v3 modeling system modified for long climate simulations at ANL
  - Performed comprehensive studies of resolution dependence on the simulation of extreme meteorological events such as the 'Perfect Storm' and 'Hurricane Huron'
  - Contributed to the Program to Intercompare Regional Climate Simulations Experiments 1-B (1993 Flood) & 1-C a 15 year simulation over US
- Access Grid
  - CoG Kits are used to enable Grid-based next generation collaborative sessions (Rick Stevens)
- OGSA
  - Uses CoG Kits to enable GSS based security (Ian Foster)

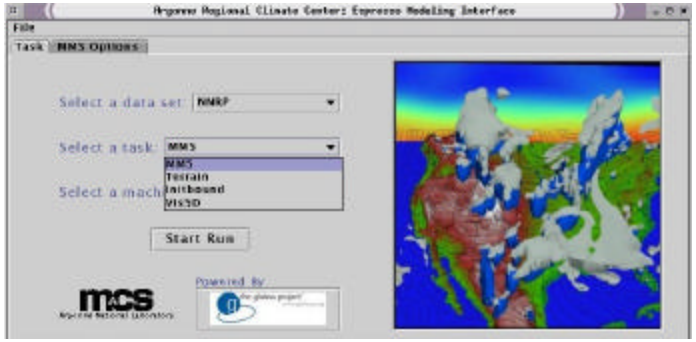

Argonne Regional Climate Center uses CoG Kits to access Grid resources

Access Grid (Rick Stevens) uses CoG Kits to enable next generation collaborative sessions

U.S. Department of Energy National Collaboratories

BERKELEY LAB

Argonne National Laboratory University of Chicago

Summary of Available Tools and Technologies

- Java
  - Globus Protocol compatible Toolkit
  - Example GUI
  - Defacto standard for portal developers
  - GSS based security for OGSA
- Third Party
  - CCAToolkit (XCAT)
  - Portal Toolkit (GPDK)
  - large user numbers
- Python
  - Globus API compatible Toolkit
  - GSS based security for OGSA
  - Example GUIs
  - Access Grid 2.0 controller based on Python CoG Kit
- Tutorials for both

Collaborations and Interactions with Other Projects

- SciDAC
  - Middleware technologies to support science portals
  - CMCS pilot
  - Middleware to support group-to-group collaboration
  - DOE Science Grid
  - Particle Physics Data Grid Collaborative Pilot
  - Earth Systems Grid II
- Others
  - Alliance Expeditions
    - Astrophysics, climatology, comp. biology, BioCore
  - ANL Midwest Reg. Climate Center
  - GridLab
  - Access Grid
  - LIGO
  - OGSA
- **There are more …**

Technology developed as part of the CoG Kit Project is today used in any project that accesses the Grid in Java or Python.

eServices Infrastructure for Collaborative Services

- Project Players
  - Kate Keahey (ANL)
  - Keith Jackson (LBNL)
  - 3 year project, started May 2002
- Project Objectives
  - Develop a unifying architecture for the Grids
    - Open Grid Services Infrastructure (OGSA)
  - Infrastructure implementations
    - Addressing the needs of scientific community
    - Language bindings: C/C++, Python
    - Performance aspects: communication protocols and encodings
  - Higher-level Services
  - Application services and infrastructure

Project Accomplishment Highlights

- Efficient implementations of OGSI
  - C implementation of OGSI (client)
    - Secure interactions
    - Fully compatible with existing implementation of OGSI
    - GT2-style wrappers for the GT3 GRAM service
    - Released in January 2003 as part of the GT3 alpha
  - Python implementation of OGSI (client)
    - Secure interactions
    - Fully compatible with existing implementation of OGSI
    - Automatic binding generation from GSDL
  - Ongoing effort in hosting environment implementation
- Application infrastructure
  - QoS-based application server (prototype stage)

Available Tools and Technologies

- C implementation of OGSI
  - Secure extensions to gSOAP
    - Both transport-level (httpg) and message-level security plugins
  - C bindings and implementation of Grid Services
  - GT2-style wrappers on top of GT3 style bindings
  - Available at www.globus.org/ogsa
- Python implementation of OGSI
  - Extensions to ZSI
    - Automatic binding generation from GSDL
      - New Python WSDL library
    - Transport-level security (httpg)
      - Implementing message-level security
  - Implementing Python hosting environment based on WebWare
  - Available at http://www-itg.lbl.gov/gtg/projects/pyOGSI/

Collaborations and Interactions

- SciDAC Security
- Applications:
  - National Fusion Collaboratory
    - QoS-based application servers
  - Access Grid
- Seeking application involvement
  - Initial infrastructure just released
  - Soliciting requirements
  - www.mcs.anl.gov/eservices

Distributed Monitoring Framework

People: Brian Tierney (PI), Dan Gunter, Jason Lee
Lawrence Berkeley National Laboratory

The goal of this project is to:

- Provide the ability to do performance analysis and fault detection in a Grid computing environment
- Improve end-to-end data throughput for data intensive applications in high-speed WAN environments
- Provide a unifying view to a wide range of sensor data, from network to host to application
  - Leading a GGF effort to define common protocols and data formats

DMF Software Components

- NetLogger
  - Easy to use instrumentation library
  - Ability to correlate data from various sources based on time
  - Easy way to collect data from multiple clients/servers reliably
  - Visualization and analysis tools
  - Relational DB-based archival tool
- pyGMA provides
  - Easy to use producer and consumer python library for constructing GGF-defined Grid Monitoring Architecture (GMA) services
- Activation Service provides
  - Ability to remotely trigger and collect monitoring data in running Grid applications

Major Accomplishments

- NetLogger Improvements
  - Binary format
  - Reliability
  - Trigger mechanism
  - Netarchd: mySQL-based archive for NetLogger data
- pyGMA
  - Released v1.0
  - Activation Service
  - SC02 Demo with Atlas software
- Standardization efforts
  - PPDG
  - Grid Laboratory Universal Environment (GLUE) Schema
  - Global Grid Forum (GGF) work
    - Discovery and Monitoring Event Descriptions WG
    - Network Measurement WG

Interactions with other projects

- Particle Physics Data Grid (PPDG) and EU DataGrid
  - Working closely with PPDG and EDG, and GLUE to define monitoring schemas and storage services (e.g.: Shoshani's SRM)
  - Working closely with Athena developers to instrument and monitoring Athena-based Grid jobs
- Globus
  - Instrumentation of Globus with NetLogger, and working together to design GT3 monitoring services
- Internet 2 End-to-End Performance initiative and SLAC IEPM Project
  - Defining schemas for network monitoring data
- Scalable and Secure Peer-to-Peer Information Sharing Project
  - Experimenting with using this for GMA producer discovery
- DOE Science Grid
  - Using NetLogger and pyGMA with various projects that are using the DOE Science Grid

SciDAC Network Research

Thomas Ndousse, DOE
High-Performance Networks Program

- Optimizing Performance and Enhancing Functionality of Distributed Applications using Logistical Networking
- INCITE – Edge-based Traffic Processing for High-Performance Networks
- Bandwidth Estimation
- Security and Policy for Group Collaborations

Optimizing Performance and Enhancing Functionality of Distributed Applications using Logistical Networking

Principle Investigators

- Micah Beck, UT
- Jack Dongarra, UT
- James Plank, UT
- Rich Wolski, UCSB

Project Objectives

- Developing advanced network and middleware services for remote storage management, making Logistical Networking fast, robust, scalable, and easy to use;
- Experimenting with and exploring the capabilities of this middleware in a range of innovative SciDAC applications with high performance and/or unique functionality; and
- Disseminating these technologies and applications to the SciDAC community on a special testbed made available for this purpose.

Logistical Networking Accomplishments

- Publications
  - 7 Publications in peer-reviewed conferences& journals, including
    - Sigcomm 2002
    - CCGrid 2002
    - Future Generation Computer Systems
    - IEEE Internet Computing
    - ACM/IEEE Transactions on Networking (submitted)
  - 5 UT Technical Reports published
- Overlay Routing and Overlay Multicast using IBP Demonstrated at
    - Supercomputing 2002
    - iGRID 2002
- 180 Mbps Transfer Performance
    - 180 Mbps UT to ORNL
    - 80 Mbps NC to ORNL

Software Deliverables

LoRS Tools – v 0.80 Logistical Runtime System Tools	**IBP – v 1.2.1** Internet Backplane Protocol
▶ Overlay routing ▶ Fragmentation and striping ▶ High performance multi-threaded transfers	▶ supports TCP, UDP, and multicast Data Movers ▶ IPv6 compatibility ▶ RAM depot
L-Bone – v 1.0 Logistical Backbone	**ROLFS – (Spring 2003)** Read Only Logistical File System
▶ Find IBP depots by network proximity ▶ Find IBP depots by geographical proximity ▶ Monitor inter-depot throughput	▶ Stores and manages exnodes ▶ Data migration for fault-tolerance ▶ Unix-like file system semantics

SciDAC Collaborators

Terascale Supernova Initiative ▶ Anthony Mezzacappa, ORNL ▶ John Blondin, NCSU	**Lawrence Berkeley National Laboratory** ▶ Arie Shoshoni
Oak Ridge National Laboratory ▶ Randy Burris ▶ Nageswara Rao ▶ Tom Dunnigan	**Jefferson Laboratory** ▶ Chip Watson

INCITE – Edge-based Traffic Processing for High-Performance Networks

R. Baraniuk, E. Knightly, R. Nowak, R. Riedi
Rice University

Les Cottrell
SLAC

Wu-chun Feng
LANL

INCITE – InterNet Control and Inference Tools at the Edge

- Goal: transform modern high-speed inter-networks into manageable and predictable systems
  - new theory and methods for network monitoring, probing, and analysis based solely on edge-based measurement at hosts and/or edge routers
  - enable network-aware and resource-aware OS's and applications

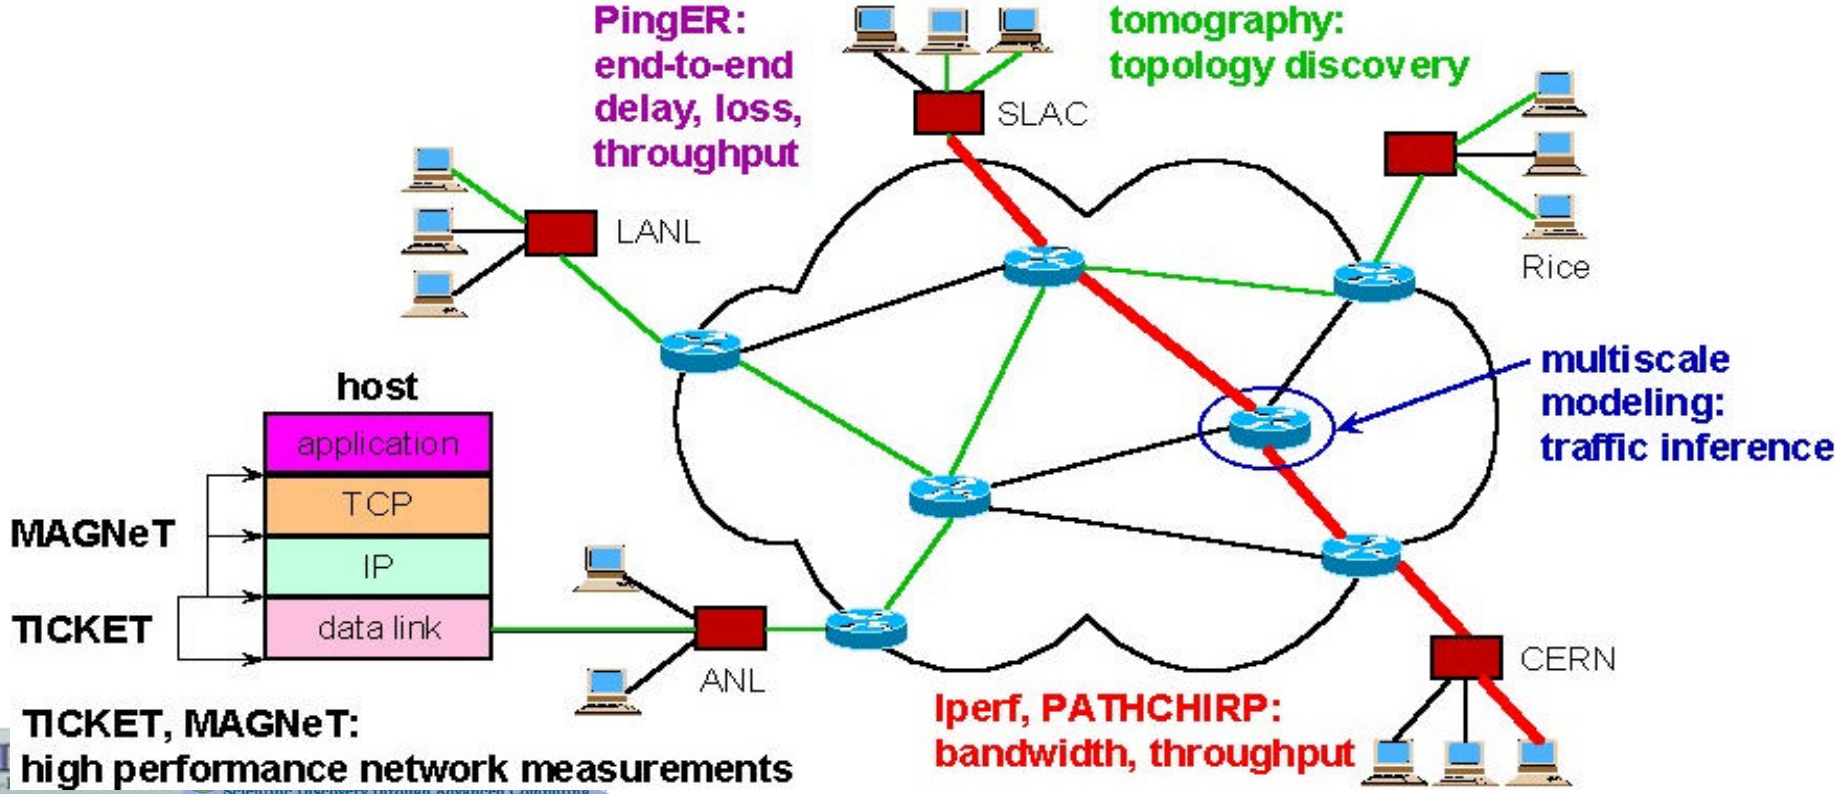

INCITE Tools and Technologies

- Monitoring tools
  - software "network oscilloscope" for clusters and grids

MAGNET	Monitoring Apparatus for General kerNel- Event Tracing
MUSE	MAGNET User-Space Environment
TICKET	Traffic Information-Collecting Kernel with Exact Timing

- Edge-based probing tools
  - inject probe packets to determine network conditions and characteristics

pathChirp	efficient estimation of available bandwidth and delay on end-to-end paths
netTomo	localizes delays and losses on individual links
netTopo	discovers internal topology from the edge
PingER	ping end-to-end reporting

- Statistical analysis and modeling tools

Current (& Pending) INCITE Users

- Particle Physics Data Grid Collaboratory Pilot
- SciDAC Center for Supernova Research
- Scientific Workspaces of the Future
- San Diego Supercomputing Center
- Transpac at Indiana University
- Rice Teraflop Cluster
- ns-2 project
- Telecordia
- TeraGrid
- CAIDA

incite.rice.edu

DOE Program Manager: Dr. Thomas Ndousse
High-performance Networks Research
Phone: 301-903-9960
tndousse@er.doe.gov

End-to-End Path Modeling

- Temporal localization of losses, delays
- Probing tools:
  - pathChirp (available bandwidth)
  - Delphi (delay distribution, cross-traffic volume)

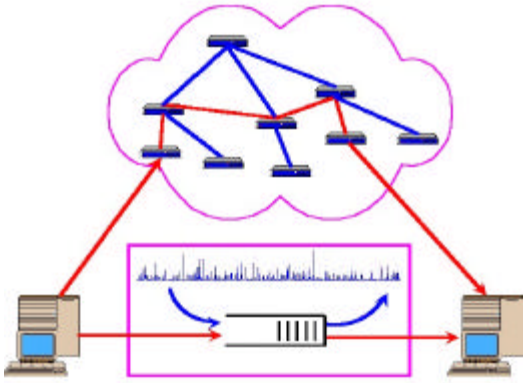

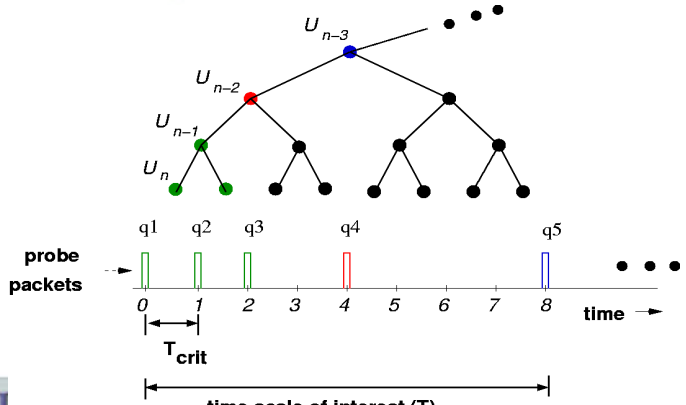

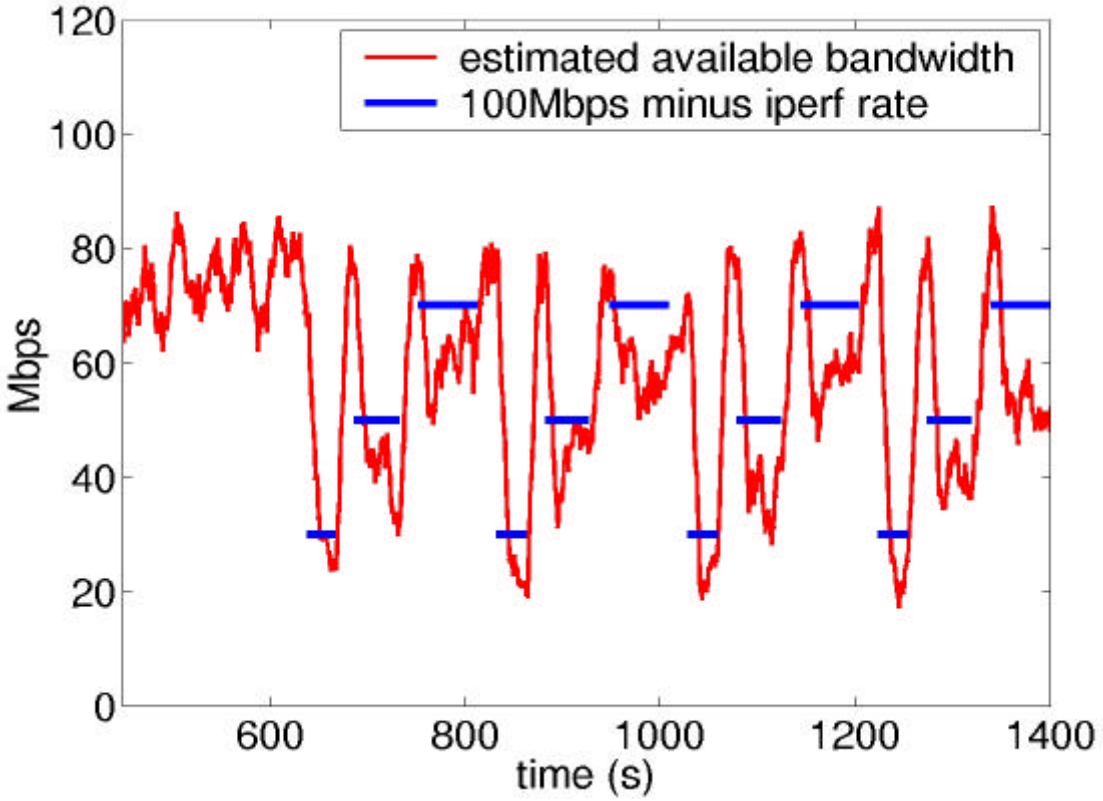

Network Tomography

- Spatial localization of losses, delays
- Probing tools:
  - neTomo (localize delay, loss)
  - neTopo (discover network topology)

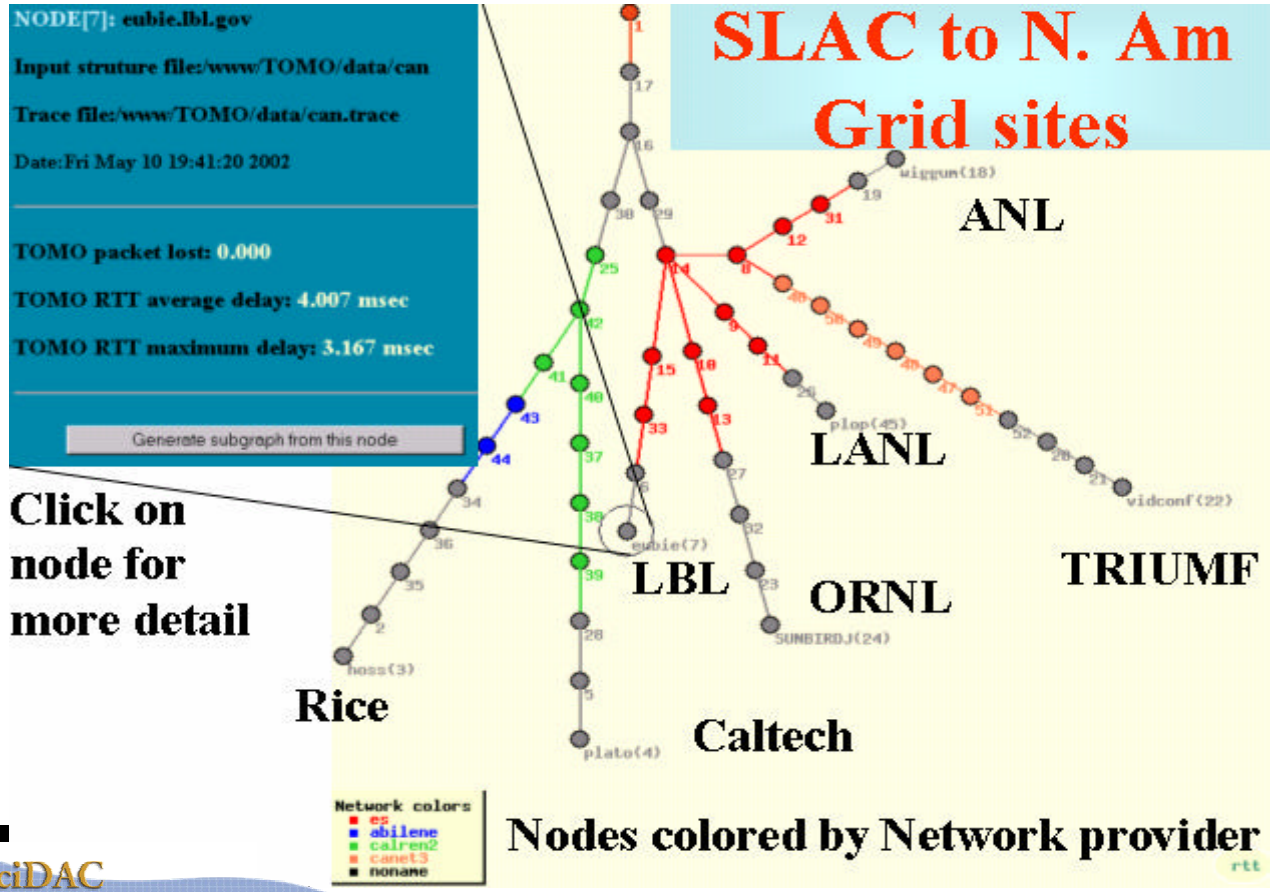

MAGNET and MUSE

- Software oscilloscope for clusters and grids

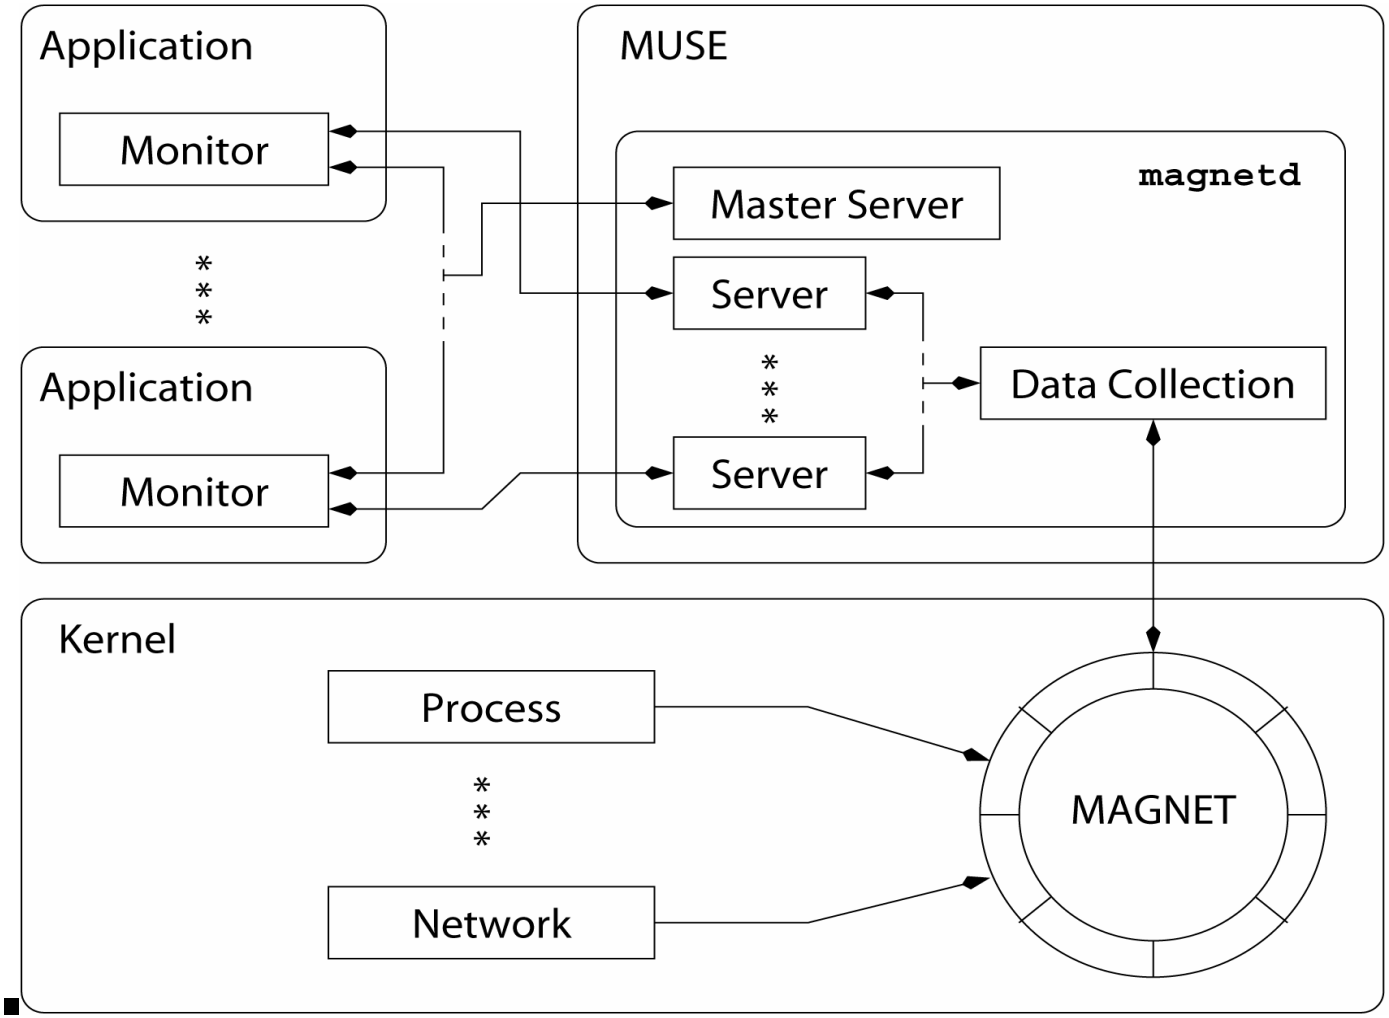

TICKET: Traffic Information-Collecting Kernel with Exact Timing

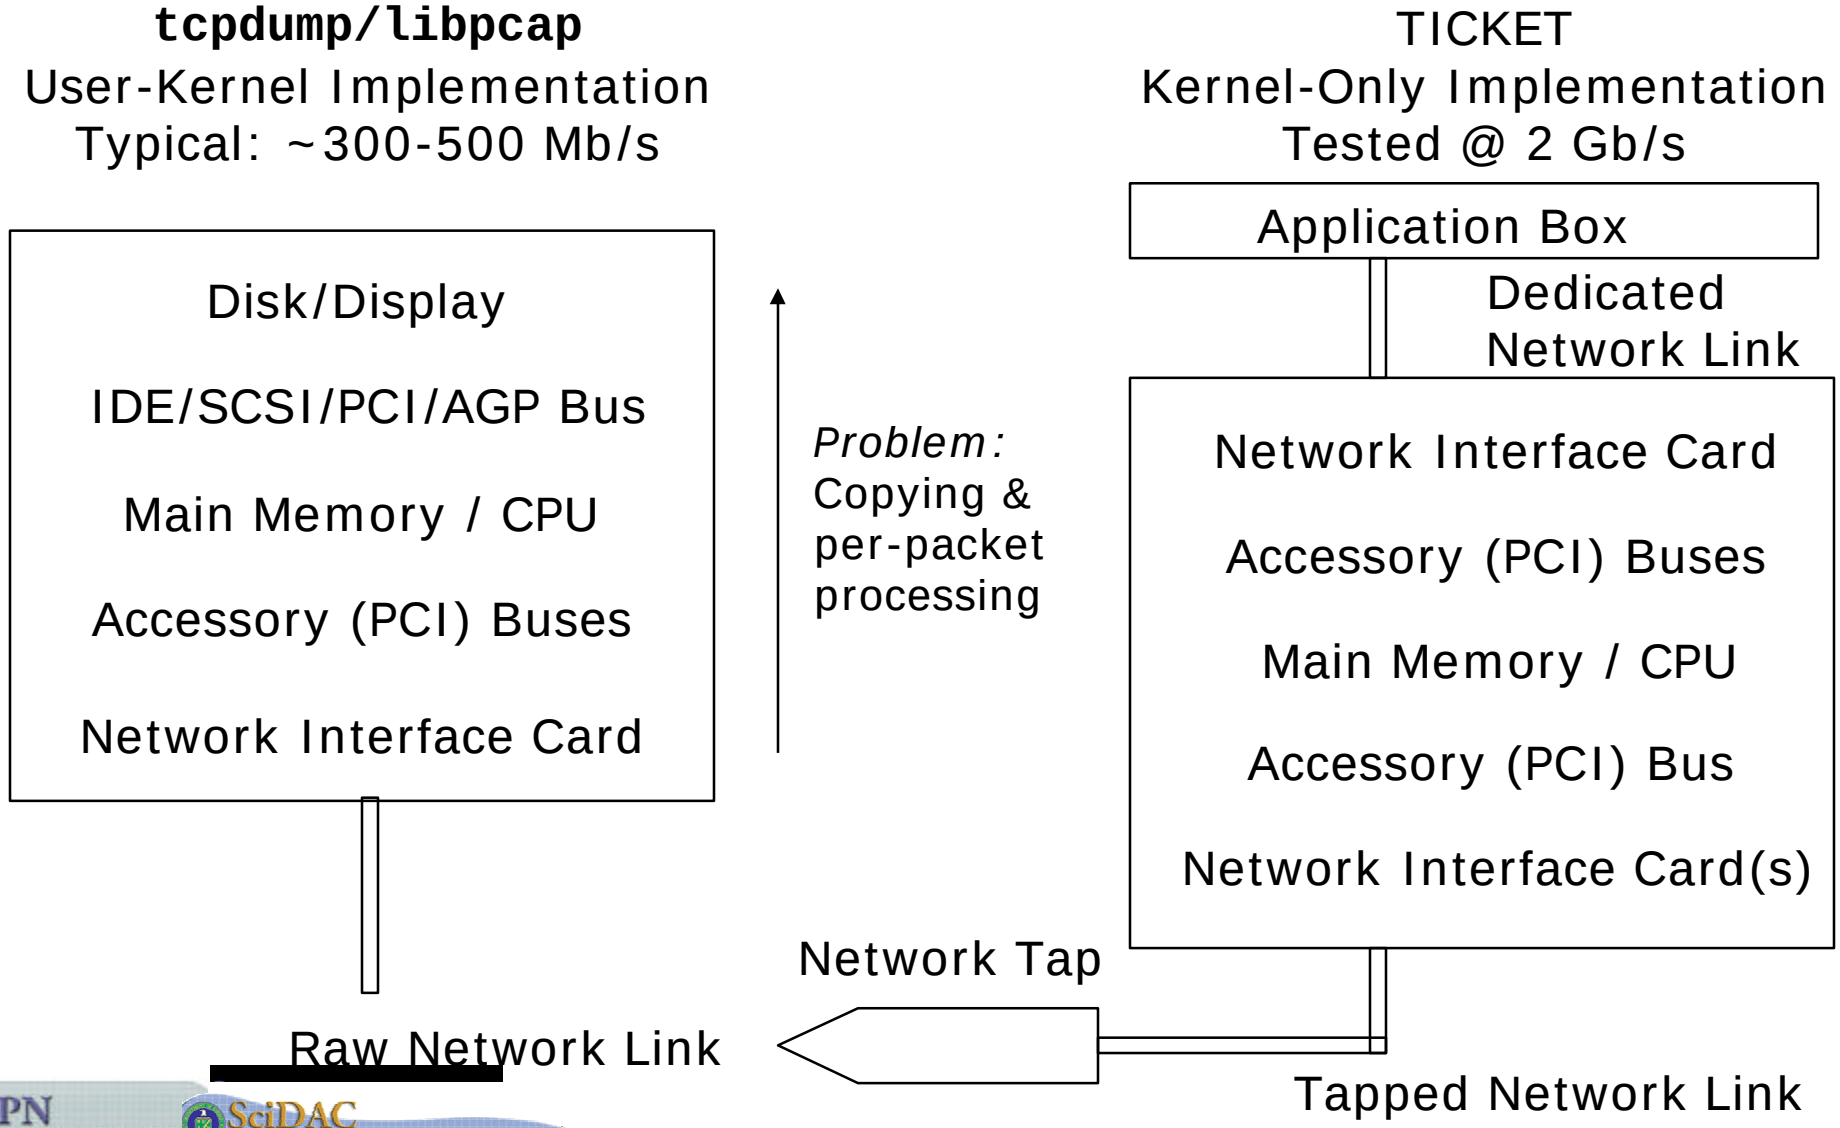

bandwidth estimation (bwest)

- CAIDA
  - k claffy
  - Margaret Murray
  - Nevil Brownlee
- Georgia Tech
  - Constantinos Dovrolis
  - Ravi Prasad
  - Manish Jain

Goals

- **Improve existing techniques and tools.**
- **Test and integrate tools into DOE and other network infrastructures.**
- **Incorporate bandwidth measurement methodologies into applications or operating systems.**
- **Examine ways in which routing protocols, traffic engineering, and network resource management systems can use accurate bandwidth estimation techniques to improve overall network efficiency.**

Results of bwest tool testing at high bandwidth

- layer 2 store and forward devices foil measurements by some bwest tools
  - ***clink, pathchar, pchar*** don't work
- latest versions of ***pathload*** and ***pathrate*** measure capacity and available bandwidth up to OC-12 (640M) rates.
- ***pipechar*** doesn't measure available bandwidth.
- GigEther paths are particularly challenging because network interfaces batch interrupts.

Available tools and technologies

- **pathload** measures e2e available bandwidth.
- **pathrate** measures e2e capacity.
- **CalNGI Reference Test Lab** offers tools comparison capability.
- **ANEMOS** GUI coordinates ***ping*** and ***pathload*** (or other tools') measurements of delay/loss/availbw on multiple paths; displays results with MRTG graphs; archives results in ***mysql*** database.

Benefits to SciDAC community

- new tools ***pathrate*** and ***pathload*** work on paths with layer 2 store and forward devices
- new application technique for automatic socket buffer sizing does not require TCP or OS modification
- bwest workshop planned for summer 2003;
  seeking wide ranging participation from DOE groups.

SciDAC Security and Policy for Group Collaborations

PIs: Steve Tuecke (tuecke@mcs.anl.gov),
Carl Kesselman, Miron Livny
http://www.mcs.anl.gov/dsl/scidac/security/

- Science today is less about individuals working in isolation and more about large, distributed, dynamic collaborations
  - ESG, EDG, CMS, PPDG, GriPhyN, NEES, etc.
- Goal: Research and develop tools for establishing and maintaining structure of distributed collaborations

Challenges

- Creation and management of large distributed collaborations has a number of challenges
- Collaborations are often dynamic in both participants, resources and sites
- Must form overlay of collaboration and local policies
  - Local policies must remain in effect while allowing collaboration policies to overlay as allowed.
- Integration of local security mechanisms (Kerberos, AFS, etc.)
  - Strive for single-sign on to keep user sanity intact

Community Authorization Service

http://www.globus.org/Security/cas/

- Enables community policy to be overlapped over local domains of collaboration
  - Outsource policy admin to VO sub-domain
  - Enables fine-grained policy
- Resource owner sets course-grained policy rules for foreign domain on “CAS-identity”
- CAS sets policy rules for community users
  - Requestors obtain capabilities from their local CAS that get enforced at the resource
- Integrated with SciDAC Earth Systems Grid software
  - Being evaluated by others (PPDG, NFC, DOESG)

Key Collaborations

- CAS prototype based on Python CoG Kit from SciDAC Cog Kit Project
  - Early use of Python Cog Kit, provided feedback to developers.
- Earth Systems Grid
  - Continuing integration with ESG software
  - Demonstrations at SC 2001 and SC 2002
- PPDG, DOESG, Fusion Collaboratory
  - Collaborations for requirement gather, feedback effecting releases, plans for deployment